COMPARING INTERACTION TECHNIQUES IN A VIRTUAL REALITY MUSEUM FRAMEWORK

USING PRESENT-DAY TECHNIQUES TO ACCESS THE PAST

Master's Thesis by Dirk Verhagen

Department of Mathematics and Computer Science at the Eindhoven University of Technology, the Netherlands
Project performed at the Re-Flex flexibility learning centre at Lund University, Sweden

Supervisors

Prof. Dr. P. De Bra (Eindhoven University of Technology)

Prof. Dr. K. Tollmar (Lund University)

TABLE OF CONTENTS

1 ABSTRACT 1

2 INTRODUCTION 2

2.1 OBJECTIVES 2
2.2 A VIRTUAL REALITY MUSEUM 3
2.3 OVERVIEW 4

3 PROBLEM DOMAIN & DESCRIPTION 5

3.1 VIRTUAL REALITY DOMAIN 5
3.1.1 Terminology 7
3.2 PROBLEM DESCRIPTION 8
3.2.1 Target group 8
3.2.2 Choice of interaction device 8
3.2.3 Tasks in a virtual museum 9
3.2.4 Design of the application 9
3.2.5 Evaluation of the design 10

4 PREVIOUS AND RELATED WORK 11

4.1 INTRODUCTION 11
4.2 CULTURAL HERITAGE IN VR 11
4.3 INTERACTION DEVICES 13
4.3.1 Conventional 3D input devices 13
4.3.2 The WiiMote 14
4.3.3 Considered Interaction Devices 16
4.4 USABILITY IN VR 18
4.4.1 Differences with traditional usability and problems 18
4.4.2 Evaluation Methods 20
4.5 DESIGN OF 3D INTERFACES 23
4.5.1 Top-Down Design 23
4.5.2 Bottom-up building blocks 25
4.6 OTHER ISSUES 27
4.6.1 Something about data structures 28
4.6.2 The concept of Edutainment 30
4.7 FURTHER READING 32
4.8 SUMMARY 33

5 DESIGN OF THE VIRTUAL MUSEUM APPLICATION 34

5.1 INITIAL USER STUDY & TASK ANALYSIS 34
5.1.1 Method 34
5.1.2 Result 36
5.1.3 Conclusions 40
5.2 USER TASK SCENARIOS 41
5.2.1 Considered Tasks 41
5.2.2 Setup of the Scenarios 42

5.2.3 Scenario One: Taking a stroll around the museum 42
5.2.4 Scenario Two: Exploring content related to an exhibit 43
5.2.5 Scenario Three: Playing a small game 43
5.2.6 Scenario Four: Immersing yourself in the 'Historical Simulation' 44
5.3 The Design of the Virtual Museum Framework 44
5.3.1 Environment 44
5.3.2 Overall look and feel 45
5.3.3 Navigation 45
5.3.4 Selection 47
5.3.5 Mapping the input device 50
5.4 Overview 55
6 THE APPLICATION 57
6.1 Software Used 57
6.2 3D Models 58
6.3 A Technical Model of the Application 60
6.4 From Design to Application 62
6.5 Using the Interface Devices 63
6.6 Adapting to different displays 65
7 USER STUDY 66
7.1 Objective 66
7.2 Hardware Setup 67
7.2.1 Test registration 68
7.3 Test Group 68
7.4 Task Set used 68
7.4.1 Task Set 69
7.5 Variables and metrics 70
7.6 Post-test questionnaire and interview 72
8 RESULTS 74
8.1 Test group and test order 74
8.2 Metrics explained 74
8.3 Device results 75
8.3.1 WiiMote vs. SpaceBall 75
8.3.2 On Screen Hints vs. No On Screen Hints 76
8.3.3 WiiMote vs. SpaceBall when using On Screen Hints 76
8.3.4 WiiMote vs. SpaceBall without On Screen Hints 77
8.3.5 First round of tasks vs. Second round of tasks 78
8.4 Questionnaire results 79
8.5 Interview results 81
9 CONCLUSION, DISCUSSION AND FURTHER WORK 83
9.1 Interaction Devices: SpaceBall vs. WiiMote 83
9.2 On Screen Hints and their influence 84
9.3 VR Museum Framework design 86

9.3.1 Expectations vs. Possibilities *86*
9.3.2 What worked *87*
9.3.3 What didn't work *87*
9.4 POSSIBLE IMPROVEMENTS 88
9.4.1 Interface suggestions *88*
9.4.2 Device suggestions *90*
9.5 RELATED RESULTS 90
9.6 FURTHER WORK 91
10 ACKNOWLEDGEMENTS 93
A. APPENDIX: CITED WORKS 95
B. APPENDIX: USER STUDY – QUESTIONNAIRE 97
C. APPENDIX: USER STUDY – INTERVIEW 101
D. APPENDIX: FULL TABLE OF TASK PERFORMANCE RESULTS 103
E. APPENDIX: FULL TABLE OF QUESTIONNAIRE RESULTS 106

1 ABSTRACT

This project will focus on the comparison of a number of different interaction techniques that can be used in a Virtual Reality Museum framework. The aim of the Virtual Reality museum framework is to provide the user with a stimulating experience through comfortable navigation, relevant content and an entertaining educational aspect.

Central to this project is the user experience. Because of this users were involved from the beginning. The devices to be tested were selected in deliberation with potential users. A user task analysis was performed to establish the user requirements.

Using this feedback a VR framework has been created which reflects the wishes of users in the best way possible, which will allow the user to visit a virtual museum and perform tasks considered relevant in a virtual museum using the WiiMote or the SpaceBall. These two devices are evaluated and compared to each other with regards to performance based metrics. The influence of on screen hints on both the performance of the devices and on user-immersion was also measured. Furthermore the users have been interviewed and filled out questionnaires to find out more concerning their wishes with regards to possible features of a full virtual museum.

Users placed much importance on navigation and selection tasks, where manipulation was deemed less important. The easy retrieval of different types of media using the custom developed 'selection wheel' and easy travelling using automated techniques such as the ZoomBack technique were most appreciated, where as an educational game, while still very much appreciated, was still less important than navigating and accessing content. All tasks were considered very easy to learn.

It is clear that for navigation based task the WiiMote outperforms the SpaceBall. However for menu operation and selection the SpaceBall was much better. On screen hints supported both tasks quite well regardless of the interaction device and improved performance significantly, though only for tasks considered problematic to begin with, without hindering the immersion of the user.

2 INTRODUCTION

Researching cultural heritage in VR is not very new. Different efforts have already been made in the past by a number of researchers. A comprehensive overview of this can be found in (Kim, Kesavadas, & Paley, 2006). However these efforts have mostly focused on realistic modeling, enhancing sense of presence in historic simulation or the transfer of information through lifelike animations. Less attention has been given to the easy retrieval of more 'museum-like' information such as texts, videos, pictures, etc. Furthermore these applications were usually created with a very specific background or with a very specific group of people in mind and without paying too much attention to the specific interaction metaphors and devices used.

One would wish that there is a way to interact more with the environment instead of admiring how nice and life-like it looks, both in virtual and real museums. This would certainly add flavour to the museum experience as usually a normal museum visit allows for very little interaction with the environment apart from some educational hypermedia devices. Also it could allow for personalization of a museum, adding or removing objects, pathways, lighting, historical accurate simulations, etc. One could think of countless examples how this could enhance a virtual visit to your favourite museum.

During this project we have tried to create a framework that addresses some of these issues. Using and combining several proven technologies an application framework has been created where people can walk through a museum that is 'alive' and offers information-rich interaction. To create a more realistic sense of a museum we collaborated with the Kulturen museum in Lund for inspiration on our small test-exhibition. In the future this framework could be a valuable addition to any real-world museums or even one that could replace a normal museum visit. We have focused our efforts on making sure this framework challenged the user to explore information and through active use gain a deeper understanding of the material offered to the user while remaining easily navigable and usable by a wide range of users.

2.1 OBJECTIVES

The main focus of this project is on the interaction devices, their use and how they compare to each other. Eventually recommendations are made concerning specific devices on how suitable they are for specific tasks considered relevant to a virtual museum as well as recommendations for the implementation of an interface. To test these devices a framework of a VR museum is created to gain results relevant to tasks in these environments. The central research question could be stated as:

How do different devices and modes of interaction compare when used in a virtual reality museum framework?

A question derived from this is:

What sort of interaction possibilities should a virtual reality museum framework offer to provide a user with a fun and educational experience?

In the end, how 'natural' a device is functioning is perceived depends on the user, which is why extensive usability testing for the described solutions is done at the end of the project to test certain hypotheses and to provide the recommendations with some statistical substance. One of the test variables is obviously which interaction device is preferred. However since there is not always a clear ranking in these devices (Bowman, Kruijff, LaViola, & Poupyrev, 2004) the mode of interaction should also be tested. For this project we have chosen to test the use of

on-screen hints for each device since these should also have an obvious influence on task performance. In the end the question whether the choice of interaction device is very important, or that rather, the implementation and helpfulness of the interface in support of the tasks that can be done in a virtual museum is more important, should be answered. Furthermore the report should touch upon how content users are with the currently offered framework, as well as offer design suggestions for future improvements, as it seems unlikely that the first version of this framework will be perfect.

Besides the interaction devices, this project touches upon some other issues concerning a virtual reality museum which should at least be researched, after which recommendations for further research or recommendations concerning their use in a (future) complete virtual reality museum can be made.

Another important aspect of the research question is the definition of Virtual Reality Museum. To provide the reader with a mental framework we will now try and give an idea of the concept of a VR museum as used during this project, so that the reader may keep this in mind while reading the remainder of the thesis.

2.2 A Virtual Reality Museum

A Virtual Reality Museum as we saw it during this project is an environment, resembling that of a real museum, where observation of the exhibits and simple yet educational information retrieval on these exhibits is of paramount importance. Due to the concept of Virtual Reality it should be interwoven with digital sources of information somehow as well as showing the relation of this information to said exhibits. This definition did evolve during the project into more specific parts, especially those pertaining the retrieval of information.

The resemblance to a *real* museum was kept in mind to keep this virtual museum quite general. Of course one could have chosen many alternatives, such as a virtual tour through space, or an environment reconstructing certain historic events. However these are all very much related to the subject of a museum and since we were focusing on a general setup our mental idea was that of a room with any kind of exhibits, as is often seen in many types of museums.

Our final implementation is still only a framework which we designed to support our tests hence the level of detail and size of the simulation were kept relatively low.

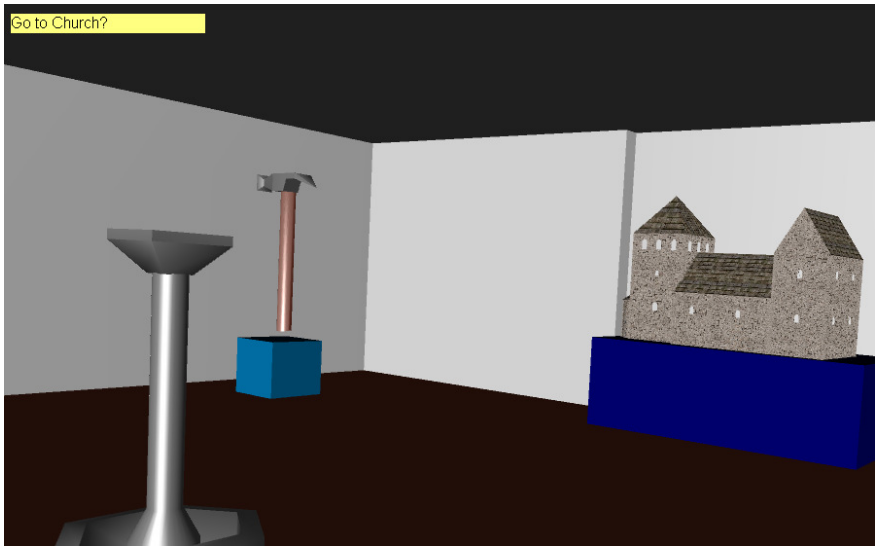

FIGURE 1: A SCREENSHOT OF THE VR MUSEUM AS IT WAS IMPLEMENTED FOR THIS PROJECT

2.3 OVERVIEW

We will start by providing readers new to Virtual Reality with a short overview of the status and definition of the Virtual Reality domain in chapter three. Using this short explanation and the concepts explained about Virtual Reality the problems that are likely to be encountered in a project such as this are described here in more detail. To provide the reader with some more background on some problems and to show that some solutions to these problems already exist we will present some work that has been done by others in the past and which is related to this project.

We will then move on to the second part of this thesis, which is our approach taken to answer the objectives set based on the research done in the previous part. We will describe the design of the VR museum framework in chapter five, commenting on why we took certain decisions and how we arrived at a better idea of what this framework might entail. This design has been made into an application which is described in more technical detail in chapter six. If the reader is not interested in the implementation he might want to skip this chapter. Chapter seven then describes the design and set-up of the extensive usability test that was done to test our design described in chapter five.

The final part of this thesis is comprised of the results and the conclusions one might make looking at the results obtained from this test. Chapter eight will contain a summarized version of statistical results comparing several modes of interaction to others based on some metrics as well as some more qualitative results. Chapter nine will deal with conclusions that might be made from these results as well as some design suggestions to improve on certain issues.

Finally, chapter ten contains a word of thanks to all the people who helped make this happen.

3 PROBLEM DOMAIN & DESCRIPTION

In this chapter we will briefly describe the Virtual Reality domain (though it is by no means a complete overview, this can be found in the references) after which we will try and provide the reader with detailed descriptions on the problems encountered. Readers already familiar with Virtual Reality will probably not find anything interesting or new in the first paragraph hence they may choose to skip it. It will not affect the understanding of this project or thesis in any way.

3.1 VIRTUAL REALITY DOMAIN

Virtual Reality (VR) is a complex field spanning many disciplines of science including, but not limited to: interaction design, computer science, psychology, graphics processing, electrical engineering, etc. All these sciences are used together to construct a *virtual environment (VE),* which is a world that usually is meant to be as real as we can make it. This VE's can be used to educate, train, entertain, inspire and simulate amongst other things (Kay M. Stanney, 2002). Virtual Reality is thus often to be experienced as real as possible. A definition of VR that extends on this point is made by Bryson (Bryson, 1994):

Virtual Reality is the use of computer technology to create the effect of an interactive 3D world in which the objects have a sense of spatial presence.

What is important here is the mention of the words 'spatial presence'. Many papers and books make a mention of 'presence' (the leading journal of MIT Press is even named after it, Presence: Teleoperators and Virtual Environments). Presence is a purely experienced phenomenon. A user is only as 'present' in a virtual application as he feels himself to be. Sense of presence is primarily achieved by presentation (Verheijen, 2004), the belief that the object the user sees are actually there. Interaction with these virtual objects and its feedback enhance this sense of presence, and therefore interaction design is an important aspect of Virtual Reality.

Through movies like *TRON* and *The Matrix* and appearances of VE's in numerous books and series (*Disclosure, Star Trek*) VR today is raised to the level of pop iconography. However when we look at the current status of VR in our everyday lives we discover that it hardly exists, with the notable exception of desktop VR which is often used for entertainment purposes (games) or in industrial applications (for example remote controlling mining operations deep underground which is done in Kiruna, Sweden).

If we take a look at the status of VR at this moment and what is shown in these movies and books there usually is a big gap. This gap is caused both by technological shortcomings as well as by the fact that some problems are simply hard to solve and still require extensive research. In the *Handbook of Virtual Environments* an agenda set by Durlach and Mavor in 1995 is presented that lists the problems that need to be solved to reach a completely immersive and real VE and their status at the time of printing of the book (2001). The only fields in which substantial advancement has been made are related to the fidelity and quality of the graphic, data access speeds and speech recognition, advancements that are usually made for other purposes than Virtual Reality (Consumer PC's today are often required to have high graphic processing power for entertainment purposes). Major shortcomings still exist in the areas of tactile feedback, real time concerns regarding interface responsiveness and graceful degradation of the rendered environment, olfactory stimulation devices, networked Virtual Environments and generalized usability studies

As has been noted many problems nowadays still exist with VR, however let's take a look at what the domain looks like today. While immersive environments such as head mounted displays (HMD's) and CAVE systems have not

found their ways into our homes yet, they are being used at universities and research laboratories, such as the one where this project was done. These display systems offer 'real' 3D worlds through the use of stereoscopic vision (giving each eye a different image so that we may see depth in the perceived world).

FIGURE 2 - A HEAD MOUNTED DISPLAY

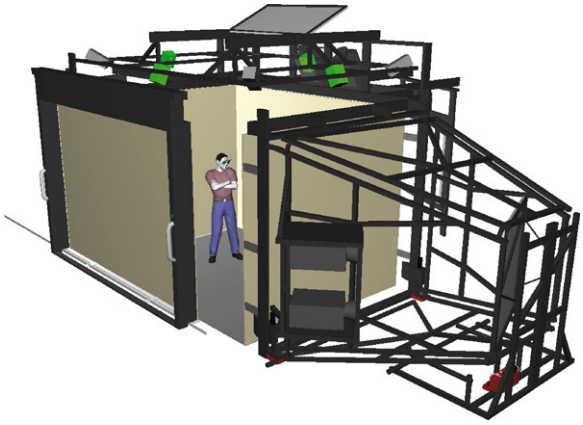

FIGURE 3 - A CAVE SYSTEM

These display systems are often used for research using simulations which require immersion and a large sense of presence. Examples of these kinds of simulations are: outer space simulations, a learning experience on how to take blood from a patient (without actually requiring human test subjects who might not be so happy at this prospect), overcoming fear of heights, a human stress test, etc. Human factors are often a very important aspect of these systems; hence they are usually designed using user-centered design.

These devices offer different advantages and disadvantages compared to each other. For example, a HMD allows the user to turn 360 degrees (assuming it is a wireless HMD) while in a CAVE system this is dependent on the number of walls used. A HMD has an integrated solution for headtracking and camera positioning while a CAVE requires another device for this. Examples like these show that technology used in VR applications has serious implications for the design of the application and may have consequences for the constraints of said applications.

As might be concluded from observing these images above the traditional mouse and keyboard are no longer suitable input devices since the user is either standing, or has no view of 'real world objects'. This is why VR research also focuses on different methods of interaction using somewhat more exotic devices. Examples of these devices are SpaceBalls, 3D Mice, Head tracking systems, data gloves, etc. which all allow and actually afford movement in the three dimensional world. A somewhat more comprehensive list can be found in (Youngblut, Johnson, Nash, Wienclaw, & Will, 1996) or for more recent versions on the internet[1]. A challenge is then left not only in getting these devices to work properly (thus not having too much jitter, accurate response times, etc.) but also to create interfaces that support these new possibilities. It is here that VR research currently still has a gap. While, as described, many technological solutions for certain problems are available, the number of applications is as of yet very small, and the VR domain is very much lacking in standards, and still exploring how normal standards could be applied to VR design.

[1] http://www.hitl.washington.edu/scivw/EVE/I.D.1.a.ActiveInteraction.html as visited in February, 2008

There are also different types of Virtual Reality. If you take a look at the CAVE system one can see that this is a 6-sided cube. However there are also CAVE systems available (such as the one at the Flexible Reality Lab) that have a smaller number of walls. These are *partially immersive* systems as opposed to the *completely immersive* systems such as a HMD or a 6-sided CAVE. The observed environment however is still fully rendered by the computer. If one starts to mix the real world with virtual imagery superimposed on this (such as in Heads-Up Displays available in most airplanes) you are entering the realm of *augmented reality*. This is still a relatively new area which could have many practical applications. As for this project, it will be done in a partially immersive system available at the Flexible Reality Lab.

The quick overview provided here is not meant to give a complete overview of everything possible in the VR world but we hoped it gave the user some insight into the technology and challenges often encountered when using Virtual Reality. It was presented to give the user some more affinity with VR for a more in-depth understanding of the rest of the project.

3.1.1 TERMINOLOGY

We will give a short overview of the most important terms used in this paper. This is by no means a complete list of relevant terminology for the VR world. A more extensive list can be found on page 18 and 19 of (Kay M. Stanney, 2002).

- CAVE
  A system where the user is surrounded by projection screens on *n* (usually between 3 and 6) number of sides. Using stereoscopic projection virtual imagery is projected onto these screens which are than perceived by the user as three dimensional, often using stereoscopic glasses. It is easy to experience the world with multiple persons (collaboratively), though to create exact stereoscopic images the position of the user's head must be known.

- Data Glove
  An interface device often used in VR systems. It can sense hand gestures and flex of fingers using fiber-optic sensors.

- Degrees of Freedom (DOF)
  Often used in the definition of user interface devices. The degrees of freedom define the amount of possible movements through three dimensional space. A mouse for example has two degrees of freedom, it can move along two axes (x and y). In three dimensional space six degrees of freedom are available (movement along all three axes, rotation over all three axes). Using these six DOF any position or orientation in the virtual world can be attained.

- Head Mounted Display
  A solution to present the user with stereoscopic images by covering both eyes with a display. This can be used just to present two different images to the user but often has head tracking built in for purposes as detecting the gaze direction.

- Head Tracking
  A system used to track the position and orientation of the head which can be important for matters as gaze-steering, gaze-selection or presenting a correct image to the user if he ducks or moves. Many different solutions exist, for example using IR tracking or ultrasound tracking.

- NunChuk
  An extension to the WiiMote which adds a joystick that can be operated with the other hand as well as two more buttons and motion-sensing capabilities similar to the WiiMote.

- SpaceBall
  An interface device created for 3D environments with a ball mounted on a 'cradle'. This ball can be moved and rotated in all directions thus providing the user with 6 DOF. Often buttons are included to extend the interaction possibilities.

- Virtual Environment
  This term is often used interchangeably with Virtual Simulation. It describes the actual world in which a user is walking around; the models, lighting, animation, etc. In some cases it includes interaction, in some cases it doesn't. This is often stated beforehand or implied in the text.

- WiiMote
  An interface device developed by Nintendo for use with the Wii gaming console modeled after a TV remote. It contains accelerometers to detect gestures as well as a certain amount of buttons for more interaction.

3.2 PROBLEM DESCRIPTION

Using the descriptions above one can flesh out the problems likely to be encountered while designing a virtual reality museum framework, and specifically the interaction, somewhat more.

3.2.1 TARGET GROUP

Since user-centered design is important we will have to have a general idea of what the target group will be. Since we are creating a framework for interaction and comparing different methods it needs not to be definitive yet, but a certain idea should always be present, so that we may keep this in mind whilst designing the application.

In this case the target group is likely to be a diverse group of people. The challenge will be in the limiting of this group. Museum visitors are a widely diverse crowd, however not everyone will be interested in using a VR application for exploring the past. The challenge lies in finding out who perhaps would not be interested in this application and by elimination we are left with our target group. In our case this would probably be persons who are slightly technophobic and need to feel comfortable with exotic interfaces at home before trying it at public places. Who exactly these people are is something we will have to keep in mind whilst designing user studies for the application.

3.2.2 CHOICE OF INTERACTION DEVICE

A very nontrivial issue for constructing any virtual reality application is the choice of interaction device. While much research has been done (Mine, 1995) (Bowman, Kruijf, LaViola, & Poupyrev, 2001) (Kay M. Stanney, 2002) the main conclusion so far is that the designer must seek the device that fits the tasks within the given constraints, without any clear 'top-device' being recommended. Since this application will be designed, not only to be useable by a large and diverse group of people, but also to be affordable for a museum, a further challenge is to find a cheap solution. Furthermore it must be obvious beforehand what the device will mainly be used for. It is important to find out exactly which task is most important and in what way that task will be implemented. Another typical consideration for the choice of interaction device is the ease of learning. While an advanced gesture system with many gestures for quick access to the functionality of an application might be really efficient, it is not easy to learn.

In a virtual museum where visitors are supposed to pick up a device and be able to fully access the functionality of the application within a couple of minutes ease of learning becomes a more important factor. Of course the design of the application (3.2.4) should also support this, since any interface can be made overly complicated if designed wrong.

Essentially this means that for this application we are looking for a device that is very affordable, easy to learn and suitable for the task at hand. The problem of finding these tasks will be explained in the next paragraph.

3.2.3 TASKS IN A VIRTUAL MUSEUM

A big problem encountered in this project is what exactly constitutes 'interaction in a virtual museum'. Therefore we will have to find out what it is a user can expect to do in a Virtual Museum, and build an application that can help him in accomplishing these tasks.

First of all there is navigation. This problem is twofold, there is movement and there is wayfinding. The metaphor used for manipulating the viewpoint (moving) should be chosen according to user's wishes and what is technically feasible. Furthermore wayfinding can be quite important in big environments. However it remains to be seen how big this first version of the Virtual Museum will be therefore this task will likely be less important. Typical problems for manipulating the viewpoint include mapping of the controls to what is happening on the screen, ease of 'looking around the environment' and the problem of controlling multiple degrees of freedom at the same time (if the navigation metaphor allows you to do this). Another problem is the choice of the interaction metaphor. There are several modes of navigation, all with their advantages and disadvantages. For a virtual museum one will probably want to give the user some amount of freedom and the ability to have an easy time in looking at exhibits, therefore these will be used as arguments for choosing a metaphor in the end.

Secondly, selection can be a big part of a Virtual Museum. Typical selection tasks can include something simple as selecting an answer to a question, as browsing through a complicated navigational structure for related exhibits and selecting the one a user would like to know more about. The problem here is finding out exactly which tasks are appreciated by the user and which aren't. Another problem is the selection mechanism. There exist many different choices for this, ranging from one-dimensional menus directly adapted from 2D interfaces to custom built 3D widgets. A careful analysis of these is necessary to make an informed choice. The choice should be made as to have a consistent selection mechanism that supports the selection tasks that can be done in the virtual museum application, without requiring too much in the way of having to learn complicated interface operations.

Last, manipulation is the task that allows you to translate and rotate the coordinates of an object, and sometimes, to resize or reshape objects. It remains to be seen how one could use this in a virtual museum since it is not a classical task in a museum. This however, could be the 'extra' thing the virtual world could offer. The opinion of the user should be leading here, since it is a relatively new question, and a more complete application will probably be necessary if the user is to appreciate these tasks in the context of a complete Virtual Museum instead of simple tasks that are separately made as showcases in a real museum.

3.2.4 DESIGN OF THE APPLICATION

Since the end result should be an application that provides us with *usable* examples of tasks that can be done in a Virtual Reality Museum (as described in the objectives) a question is: How to design the application in such a way that the focus is on usability? Different approaches for this can be chosen. It should be decided if what is being made is a completely new application, or something existing for which there already are existing evaluation reports on which the approach can be chosen.

What also remains to be seen is what, if any, is the difference between design of normal software applications and designing for Virtual Reality. Virtual Reality will more likely be focused on form, working solutions and interactive systems instead of accepted software design paradigms and well grounded development frameworks. Of course there also exist development frameworks for interactive systems in normal software design hence we should take a look if those methods are applicable to Virtual Reality design.

Designing the actual interface can also have its very own unique problems. Since we are designing a three dimensional interface there are some obvious differences with designing for desktop interfaces. We will have to try and leverage the advantages offered by this third dimension somehow in our design solution, and to make sure that our final design is not just a glorified desktop application.

Then there is the matter of documenting the design. As far as we have been able to determine there is not yet a standardized way to document your VR application and matters such as the interaction in the application, the technical means (such as databases), etc. While such methods exist in traditional computer science (e.g. UML or ORM) and to some extent in usability design (Use Case diagrams, outlines, interaction schedules) an application like this combines both and might therefore need a new approach.

3.2.5 Evaluation of the design

In the end the design should be evaluated. It is stated by different authors that usability evaluation in VR has some fundamental differences with traditional usability evaluation. For example, the sense of presence as described in the beginning cannot be evaluated by traditional means, furthermore there is a third dimension here for the interface which presents new opportunities which should also be evaluated, such as making use of the sense of proprioception (the sense of how your body is positioned).

Also, there is very little usability research done yet. Most research has focused simply on what is possible, and not so much on if that which is possible is also usable (Bowman, Kruijff, LaViola, & Poupyrev, 2004). That means that methods used by others to test their applications for usability are scarce and not always right for every specific application. How to consider and evaluate multiple devices in the limited time available is certainly a challenge.

Then there is the problem of what exactly should be measured to assure that the Virtual Reality museum is good, and in the end, what variables should be measured to be able to make interesting conclusions that add something to what already has been done in the field of Virtual Reality research? Old research should not be repeated. What one would like to measure is if the virtual museum could be entertaining, educational and usable. These things are not defined in an exact manner and one of the challenges will be to design tests and questionnaires to answer these questions. Furthermore, the planned alternatives (e.g. two interaction devices, different modes of interaction through on screen hints), these alternatives should be compared in an unbiased and clear way, keeping in mind that many seemingly unrelated variables could influence the outcome of these tests and thus the conclusions.

4 PREVIOUS AND RELATED WORK

4.1 INTRODUCTION

As presented in the domain description in the previous chapter research in the field of Virtual Reality area is as of yet still incomplete and lacking in many fields. This does not mean however that there is little research done. While it may not compare to more established fields there is still a lot of research available on this topic which we can and should use to address some of the problems described in the previous chapter.

We will now present several articles or chapters from these books and indicate in what way they are useful and relevant to this project, in what way they might be useful for future extensions of this project as well as some comments on the articles themselves.

4.2 CULTURAL HERITAGE IN VR

First off we start by presenting some interesting concepts developed about Cultural Heritage, since this seems like an obvious starting point for a project like this. This is not the first project to bring Cultural Heritage to Virtual Reality worlds and make them available to the public, far from it actually. At the flexibility learning centre where this project was done, a project was already completed[2], creating a historically accurate representation of the (no longer existing) Drotten church. The report was unfortunately in Swedish but exploring this simulation provided some inspiration for what works in bringing a simulation to life, and we will see some suggestions using simulations like these and the findings presented in this thesis in the end. The models and the report are still available on the web at the URL provided in the footnote.

A paper published in *Presence* that would support the advanced modeling techniques used in the creation of the Drotten church was written by Young-Seok Kim et al. (Kim, Kesavadas, & Paley, 2006). Kim et al. describe the setup of a virtual museum of a site that doesn't exist anymore for the purpose of research and education. Researchers can experience and explore the Northwest Palace of King Ashurnasirpal II (883-859 BC) and, whilst being in the simulation, access historical records. However, these seem not to be records as we would imagine them in a virtual museum for the public, but rather tablets or pictures of artifacts that are interesting to researchers. The ordering of, or how to navigate these records is never made explicit.

The focus of this project is on the historical depiction of certain scenes in this palace and using these scenes, to teach history in a more tangible way. An example is given where Kim et al. animated the king getting up in a very slow way, thus illustrating the weight of his garments. By providing full body immersion (using a CAVE environment) they try and create a sense of presence and enhance the bond with history. A conclusion is that this bond does reinforce the learning process and involvement in the material which seems to indicate that our idea for this project has some merit.

The article by Kim et al. is related to this project in the sense that it also tries to provide an environment that can teach users something about history by immersing them in it. However the users here are explicitly stated to be researchers and archeological students, hence the interface and possible tasks in this virtual site museum are more geared toward expert users. It does however present an interesting list of already existing 'virtual heritage

[2] http://www.reflex.lth.se/culture/kulturen/ as visited in February, 2008

projects', which all seem to be geared towards accurate reconstruction of some event or place, rather than the transfer of information to a wide public in an interactive manner, on any given subject. This is where this project might provide a new insight to cultural heritage.

Furthermore this paper notes that many VR simulations need a more user-oriented approach, and then proceed to not use this approach. There is no mention of a test before the design phase and the interface seems to be using an approach with virtual artifacts using a data glove to manipulate these artifacts, something that in 2006 is hardly new technology anymore. Whether or not these manipulations actually support tasks that can users would like to see is not mentioned nor tested in this paper. The focus is more on providing an accurate environment through physics and animation. It is also here that we can learn something and try and get a better overview of the tasks that should be supported, and then design the interface in such a way that it actually supports tasks that a wide range of users consider relevant.

The implementation issues that the article by Kim et al. deals with can be important in later, more extensive uses of the end result of this project. Most importantly it deals with digitizing historical information, and models the information stream from artifact to VR simulation. It is definitely recommended reading for those who will model the next iteration of the virtual museum prototype created in this project in more detail.

The digitization of information in a virtual museum is important enough to deserve a mention to also do some research on. It is here that papers and reports, published on the CHIP website[3] have provided some insights into personalized access to cultural heritage. A report by Yiwen Wang (Wang, 2007) describes these ideas in greater detail.

Wang describes the use of an ontology to structure data in such a way to bridge the gap between the limited vocabulary of 'beginners' in the world of musea, and experts. A problem here is that while an expert can know that he likes an 'impressionist' style, a 'beginner' might not. By gaining insights on a user's personal preferences through a rating system, which is compared to that used at amazon.com by Wang, the system can make recommendations and detect that this user does in fact, like the impressionist style.

In a sense Wang tries and create a sense that the normal museum monologue becomes more of a dialogue. This is exactly what they are trying to accomplish at the moment using perceptual user interfaces. In a paper published at the Symposium on Intelligent Information Media (Turk, 1998), Matthew Turk of Microsoft says the following: "The ultimate interface is one which leverages these natural abilities, as well as our tendency to interact with technology in a social manner". Therefore the efforts by Wang could well be used in the future interface design of the virtual museum, considering it is certainly a step in the direction of perceptual interfaces. At the moment however the described model was still in demonstration and research phase and not directly useable. Furthermore its dependency on a user's rating could also prove its Achilles' heel. In the user task analysis performed for this project, users often indicated that while they would like a recommendation system, they would probably not explicitly rate exhibits that they encountered. The researchers have however taken note of this and in an earlier publication (Rutledge, Aroyo, & Stash, 2006) a more inferred approach is described that uses an engine to infer knowledge rather than explicitly asking the user. Furthermore it describes a simple system where a user simply indicates whether he liked, was neutral, or didn't like a certain exhibit, resulting in a much less cumbersome and arbitrary scale than the one used by most webstores (the '5-star' system).

[3] http://chip-project.org/ as visited during the entire project

As it stands now a simple rating system could certainly be used in the virtual museum. Inferred knowledge is even easier to gather in this environment since you don't need RFID tags or PDA's to help you navigate or gather information, the environment is already completely digital and so called 'sensors and actuators' can be used in an intelligent manner to gather personal data about the user. What should be taken care of is seperation of the presentation and the underlying data, making sure that later additions to the structuring of the historical data can be changed. This might however also mean that new views will have to be developed to support this data and navigation trough its structure, research on this is very limited in VR and the standard components (e.g. for treeview) for desktop applications are not available for VR applications. Hence the coupling between data and view is not completely as loose as in a standard MVC application. We will get back to research on the ordering of information in section 4.6.

4.3 Interaction Devices

In the field of interaction devices much research has been done in the past years. This research however, mostly concerns solutions that have been around for a while now such as data gloves, spaceballs, joysticks and the like. The newer Nintendo WiiMote which is considered for this project is unfortunately still very new. For this reason this subject was only researched in someone else's master thesis, instead of published papers.

A common complaint about this research is that much of it is device driven as opposed to user-driven. However for this section this is not so much of a drawback since we are actually interested in finding out as much as we can about the devices. How useable they are to the end-application is something which will have to be researched during this project. Since the application being created for this project is new, the application-specific usability of each device will have to be looked at closely anyway and research in this direction is not directly necessary.

A good place to get a simple overview is the World Wide Web; a URL has already been provided in the introduction (Chapter 2) however one can also take a look at Wikipedia[4] for a comprehensive list of interaction devices. However, limitations of these online articles are obvious. While these sites have pretty complete lists of devices and descriptions of what they do, research on how well they are suited is often lacking, as well as a more intricate description of their abilities and what sort of general tasks they are suited for. Even though this is not necessarily the best for our application, using accepted paradigms might prove useful in any case.

4.3.1 Conventional 3D input devices

An article which discusses the use of physical devices, and how they compare to 'virtual' devices is written by Mine (Mine, 1995), which lists the use of physical devices and concludes that while they are often readily available, they are often counter intuïtive and in this case limited. Virtual controls (e.g. a virtual joystick or steering wheel) have the drawback that they have no haptic feedback. How these virtual controls then are operated is left to the imagination of the reader, but one would imagine that some sort of interaction device is also necessary here (even if it is just a camera that tracks body motion). Mine then proceeds to give an oversight of the possibilities for movement, selection and manipulation in a 3D environment. For example, for navigation Mine points out that one can pick to use actuators to move the user's avatar, or that the user can do this his or herself. Depending on issues like these the device should be picked. For precise movements a joystick is then recommended. The paper lists a lot more considerations for selection and manipulation, but mostly they are just that, considerations. In the conclusion Mine notes that this paper has aimed to provide an overview of what is possible rather than a solid list of conlusions. However, this paper is for that very reason, a good place to start research into interaction in 3D

[4] http://en.wikipedia.org/wiki/3D_Interaction#Input_Devices as visited from February to April, 2008

worlds. Later taxonomies for input modes and devices by Bowman seem to be influenced by descriptions of concepts that are well-worded in this paper. The obvious criticism is that Mine's report is already 13 years old, and therefore somewhat dated. The list of devices available to operate a virtual environment (joysticks, dials and knobs) is very limited and the limitations of not mapping in a clear or natural way to the Virtual Environment that were noted in this paper are not neccesarily true for devices available in these more modern times. Therefore more research is needed.

A more contemporary article is written by Jesper Kjeldskov (Kjeldskov, 2001). In this article an overview is given of several interaction devices (a headtracker in different modes, a joystick, a trackball and a spacemouse) and some advantages and disadvantages when they are used in partially immersive displays (such as panorama screens) or fully immersive displays (a Head Mounted Display or six-sided cave). An important conclusion being drawn by Kjeldskov is that head tracking is not perfect for partially immersive displays (which are the ones being considered for the virtual museum, as fully immersive solutions are likely too expensive for any museum to realistically consider). Users have problems with this because the virtual world 'ends' while they can still turn their heads. It might be reasoned that the same problem happens with torso tracking.

Concerning movement Kjeldskov concludes that motion tracking (tracking the position of the body and placing it inside the virtual world in a corresponding location) whilst being useful and natural, needs the aid of another interaction device such as a joystick to move larger distances in the virtual world, regardless of the display used. A difference with Mine's article here is that here joysticks are called 'unprecise' for movement. However this may be due to the comparison to more precise devices such as a trackball and a spacemouse which are judged to be more precise here.

Furthermore, a conclusion by Kjeldskov that could be important is that non tracked interaction devices (as opposed to head tracking systems for example) work the best in a partially immersive display while working slightly less well in fully immersive displays as users can be prone to turn their bodies (which is tracked) and feel desoriented.

4.3.2 The WiiMote

Some very modern devices have very limited research available,and included in these devices is the WiiMote. A researcher who is particularly interested in using the WiiMote in many different scenarios has however made available quite a lot of material on the world wide web[5] including a list of demonstrations which show the potential of the device. This has become an inspiration for researchers around the world as the demonstrations are very impressive, and well known in the interaction design field. We will now focus somewhat more on the WiiMote as it is a relatively new device which has not been used in the research field a lot.

Another document which contains a much more detailed description of the WiiMote (in Dutch, or Flemish as the writer would probably want it cited) is the master thesis of Gilles Vermeulen (Vermeulen, 2008). This thesis also touches on some more exotic devices such as the Phantom and the Falcon (Figure 4 and 5) but they are designed with the purpose of haptic feedback in mind, which is why we shall not delve into those.

[5] http://www.cs.cmu.edu/~johnny/projects/wii/ as visited during the entire project

FIGURE 4 - THE FALCON

FIGURE 5 - THE PHANTOM

In this master thesis the WiiMote, its possibilities and its limitations are discussed. As already described in the domain description a WiiMote can sense gestures. More accurately however would be to say it senses acceleration. For this it uses accelerometers which unfortunately have a limitation. They are unable to cancel out the force of gravity, nor are they able to detect slight tremors. A more detailed explanation is available on the internet[6]. This means that the WiiMote will not have 100% accurate gesture recognition. At the moment an addon using gyroscopes instead of accelerometers is announced by Nintendo which will make this possible but it is unknown when this will be available so it is ignored for this project. To understand what exactly *is* possible with the WiiMote we will take a closer look at what is inside this little device as is described in detail by Vermeulen.

As already mentioned the WiiMote is able to sense gestures, or rather, motion and use this as input for any programs. While it cannot differentiate gravity, if the WiiMote is held immobile it does measure the acceleration of gravitational forces (9.8 m/s^2). This means that even though accurate position tracking is not possible using the accelerometers, some orientation tracking actually is. One can measure roll and tilt since these affect gravitational forces. Yaw however has no gravitational element and this is why this sideways motion is not measurable using only the accelerometers.

Luckily, the WiiMote also has a little infrared camera. Using two infrared lightsources the WiiMote is able to calculate its yaw (by looking at the displacement of the lightsources relative to eachother) and even its distance to the two light sources (assuming one is using the standard 'sensor bar' that is usually shipped with a Nintendo Wii). This can also be used for other purposes such as VR Head Tracking and finger tracking, as shown on the website by Johnny Lee. Furthermore using four lightsources and some advanced triangulation, one can even approach real 6DOF tracking[7] which can have great implications for cheap solutions for VR 3D input devices, but still needs work as the dependancy upon these four light sources does not provide a workable solution yet for cave environment or HMD's. The announced add-on using gyroscopes could solve this issue, but that is beyond the scope of this thesis.

Furthermore the WiiMote contains several buttons, which is often easier for a user to understand than using gestures for everything. Also, using the extension port provided on the WiiMote one can connect the NunChuk, a device which also has accelerometers for each axis, as well as a 2 DOF joystick. To top it off the WiiMote also contains a little vibrating engine to provide haptic feedback, as well as four LEDs and a speaker.

[6] http://www.motusbioengineering.com/sensor-comparisons-technical-note.htm as visited on June, 2008

[7] http://idav.ucdavis.edu/~okreylos/ResDev/Wiimote as visited multiple times during the project

FIGURE 6 - THE NUNCHUK

All these possibilities combines provide us with many options for one- or twohanded gestures, button input, joystick input, pointer movement (using the IR tracking) and feedback, beit haptic, audio or visual. Visual and Audio feedback however is probably best left to the VR system.

Conclusion of the research by Vermeulen was that while the WiiMote at the moment is gimmicky, there is a lot of technology put into a device that is considered very easy to understand by a lot of people. It deserves a more serious approach as a device that combines so many possibilities can be used for much, much more than just waving it around a bit. He made several demonstration applications which use the WiiMote in combination with an IR reflecting glove to use it mainly as a manipulation tool. This still leaves the navigation and selection capabilities of the WiiMote unexplored and it remains to be seen how well suited it is to these tasks.

FIGURE 7 - HOLDING THE WIIMOTE

4.3.3 Considered Interaction Devices

Using all of the sources listed above we have compiled a list of interesting devices some of which we can hopefully use to successfully make a usable VR museum application. We will now present a short summary of the interaction devices that were deemed relevant to this project based on the research done along with some arguments for the selection of the devices. Constraints that go for all of these devices is that it needs to be none too expensive, easy to use and have a certain familiarity with users.

Device	Argumentation
Body Tracking/Stereo Camera	It provides us with a very immersive way to interact with the virtual museum. It remains to be seen how well suited it is to locomotion and how important users find this; the immersive aspect however is probably second to none. Using your body to interact is natural for humans however somewhat more novel in a Virtual Environment.
DataGlove	The data glove can provide us with very precise gesture recognition for the hands and is extremely well suited for manipulation tasks. There are also techniques available for navigating an environment and selecting since extensive research has been done on the use of this device.
PDA/TouchScreen	A device that is well spread and available and very suited for presentation of (context sensitive) information that could otherwise kill the immersion. One can think of countless things to present on the PDA screen, for example a map. Immersion might suffer though, especially since the interface to navigate will probably be a bit contrived.
SpaceBall	A device that can move and roll in every three dimensional direction giving us 6 DOF in an intuitive way. It has also been researched quite a lot. It also resembles devices that most people know (most notably a trackball) and therefore might be intuitive to use.
WiiMote	The many possibilities in this device provide a nice mix of immersion while retaining the ability for accurate input. The remote-like design and the fact that it is well spread and well-known among the general population might also help to get people to use it (provided it reacts as they expect it to, a case of proper design).

TABLE 1 - A LIST OF CONSIDERED INTERACTION DEVICES

Other alternatives were mainly rejected on the basis of them being designed for a very specific purpose (for example certain walking in place techniques), being too expensive (SculpRox) or being very exotic and hard to implement (Complete bodysuits), therefore falling outside of the scope of this master's project.

4.4 USABILITY IN VR

While much usability research has been done in the past by illustrious names such as Nielsen and of course many others, most of that focuses on traditional desktop interfaces. While many of those lessons can be generalized towards virtual environments (for example the 10 usability principles by Nielsen (Preece, Rogers, & Sharp, 2002) apply to pretty much any interface) virtual environments have their own specific problems that are not included in traditional usability research. Some obvious examples are 'cyber sickness', 'comfort' and strain (an issue for users that are standing up) and of course, three dimensional interfaces. There are not many guidelines for the design of these interfaces yet since the usability of those interfaces has not been extensively tested (Bowman, Kruijff, LaViola, & Poupyrev, 2004). However, this same book by Bowman has a quite extensive chapter on usability testing and usability engineering of Virtual Reality application. It more or less advocates a usability engineering approach to any Virtual Reality application. Furthermore it states that while the design space is well explored by now, the designs proposed have not been assessed on usability, since usability was, for quite a while, a side-issue at most in VR research until it turned to designing and creating applications. We will take a closer look at usability this chapter, and present some other papers so that we may gain some insight on what is important when thinking of usability in Virtual Environments, and how to incorporate this in this project.

4.4.1 DIFFERENCES WITH TRADITIONAL USABILITY AND PROBLEMS

A paper that delves somewhat deeper into the differences between traditional HCI methods for evaluating usability and the more specific needs for VE's is written bij J.Tromp et al. (Tromp, Steed, & Wilson, 2003). Though most of her work applies to collaborative environments, research was needed on usability in VR. A conclusion of the introductory research performed by Tromp et al. is that a big focus in VR evaluation should be on human needs. Again the dialogue between man and machine is mentioned. Since we are moving into the virtual space, and users can use their bodies to interact, and often have avatars representing themselves in the virtual world, the interface is experienced more as a dialogue between man and machine. Once again mentioned by Tromp et al. are the specific challenges for navigation in a three-dimensional world and interaction.

Tromp et al. tested their own methods of evaluation through use of a disjunct project group and a group of testers. The testers (so, not the test-subjects, not the researchers for this project, but rather the people overseeing the test process) would report problems with evaluation methods to the project group which would then try and refine their methods and definitions. In a sense one could say a usability engineering approach of usability evaluation was taken here.

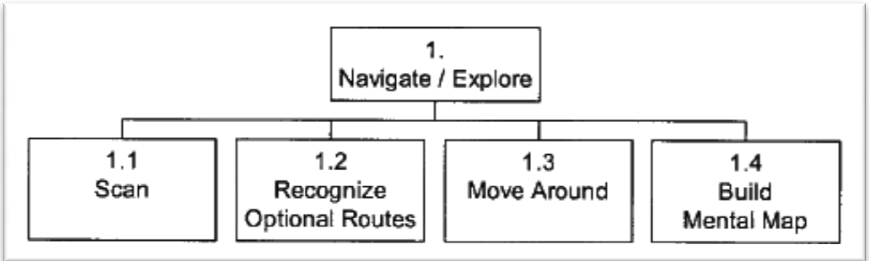

FIGURE 8 - A SMALL EXAMPLE OF A TYPICAL TASK TREE

While some problems and the evaluation of these problems were very specific to the collaborative aspect (for example,the computational load by including network trials had a definite effect on end-usability) some conclusions made by Tromp et al. are interesting. First of all a typical design of tasks called 'the task tree' was found to be insufficient. The amount of freedom experienced by test subjects and as a result the amount of actions they could take in the virtual environment dit not fit nicely into the task tree as they were proposed by researchers. A more elaborate definition of these tasks was necessary, which had to have an emphasis on important goals of the task, instead of just a definition of the task. Also a more descriptive way of the experience of the task should be included. What feelings of the user are important to take into consideration? Another

conclusion that is relevant to this project was the fact that technologies used did identify usability problems, but were not immideately focused on giving redesign suggestions. Hence problems could be identified, but what to do with those problems was a problem for the designers. A redesign was often difficult due to the facts that there are no clear guidelines and a large number of approaches can be taken in VR design, hence the redesign often suffered from new usability problems. Small redesign suggestions should therefore already be done in the evaluation section of the usability engineering process. We will see this later in chapter nine.

Bowman et al. (Bowman, Kruijff, LaViola, & Poupyrev, 2004) come up with a number of issues by indentifying problems you might encounter using usability methods in a 3D environment. Some typical problems here (and that we should take note of) are:

1. Physical Environment Issues: Physical barriers can be hard to see because the user is either wearing a HMD, or graphics are being projected on them. 3D Displays do not always allow multiple viewers, and thus observing what is happening can become a problem when testing. 3D users can be mobile, hence videotaping the user may require a wide shot, resulting in a loss of detail.
2. Evaluator Issues: An evaluator can actually disturb the sense of presence a user is experiencing. This means that the evaluator should not interfere when not necessary; as a result the virtual application must be robust and bug-free so the evaluator does not have to interfere. Also tasks should be explained in great detail before the user will start on them.
3. Hardware Issues: Hardware is often not so robust and more complex than traditional UI hardware. Some help might be required here. Furthermore, often multimodal inputs are used in 3D UI's, since each input-stream will have to be captured the challenge of videoing or recording it in some way returns again.
4. User Issues: The target population of an application is often not known due to many applications being "Solutions looking for a problem". We will have to take care to make a specific application which actually addresses real-world problems. Furthermore it can be hard to differentiate between expert and novice users, as there are not many experts in 3D UI's yet. Furthermore finding design flaws might require more than just 5 users as proposed by Nielsen (Nielsen, 1993), as often a greater statistical variance is observed due to methods that are unknown to most users.
5. Evaluation type issues: Evaluation based on guidelines (expert-evaluation) is often difficult due to the lack of guidelines for VR. Performance models are, for the same reason, also less effective. Because of the complexity of VR applications more automated measurements may also need to be taken (for example, virtual distance travelled). Another important issue here is that when performing statistical analysis, it is often difficult to know which factors have a potential impact on the result. Solutions to this may sometimes make the test either overly complex, or overly simple.
6. Other issues: 3D evaluation often operates at a lower level than standard evaluation methods. For VR there is no standard set of components available with a widely known 'look-and-feel'. 3D Interfaces must often compare low-level components such as device used, or interaction technique used. Furthermore it is tempting to over-generalize the results. Because of the complex nature of 3D interfaces anything can change here.

As can be seen there are enough of pitfalls to watch out for when looking at evaluation methods and applying them to 3D interfaces. Now that some problems have been identified as well as some differences between traditional HCI methods and VR usability evaluation we can take a look at what exactly those methods are, for VR evaluation.

4.4.2 Evaluation Methods

Bowman et al. make a small list of evaluation methods that have their roots in traditional GUI evaluation, but have been succesfully applied to 3D environments. They propose the following methods: Cognitive walkthrough, heuristic evaluation (guidelines-based expert evaluation), formative evaluation, summative evaluation, questionnaires and interviews. These are the same methods also suggested in the evaluation chapter in the handbook for virtual environments (Kay M. Stanney, 2002).

Now as to the choosing of evaluation method, Bowman et. al. have some advice for that, dependant on three things. Whether the test requires user, if the test is application-specific or generic, and if the data that is being gathered should be quantative or qualitative. The first two are obvious for this application, the third less so. A paper by Schroeder et al. (Schroeder, Heldal, & Tromp, 2006) provides us with a guideline, saying that in their experience, a qualitative and a quantitive approach actually help eachother. By doing both, the disadvantages of one method might be negated by the opportunities of the other method. This seems to make sense, especially since for this project we are also trying to gather data on the entertainment and educational aspect that cannot be measured in an exact way. However simple navigation tasks should be measurable and objectively compared. Hence in the framework proposed by Bowman we now end up in a quadrant where we do require users and perform an application specific test. A more explanatory figure of this 'evaluation method space' can be obtained from a technical report published by Virginia Tech (Bowman, Gabbard, & Hix, 2001).

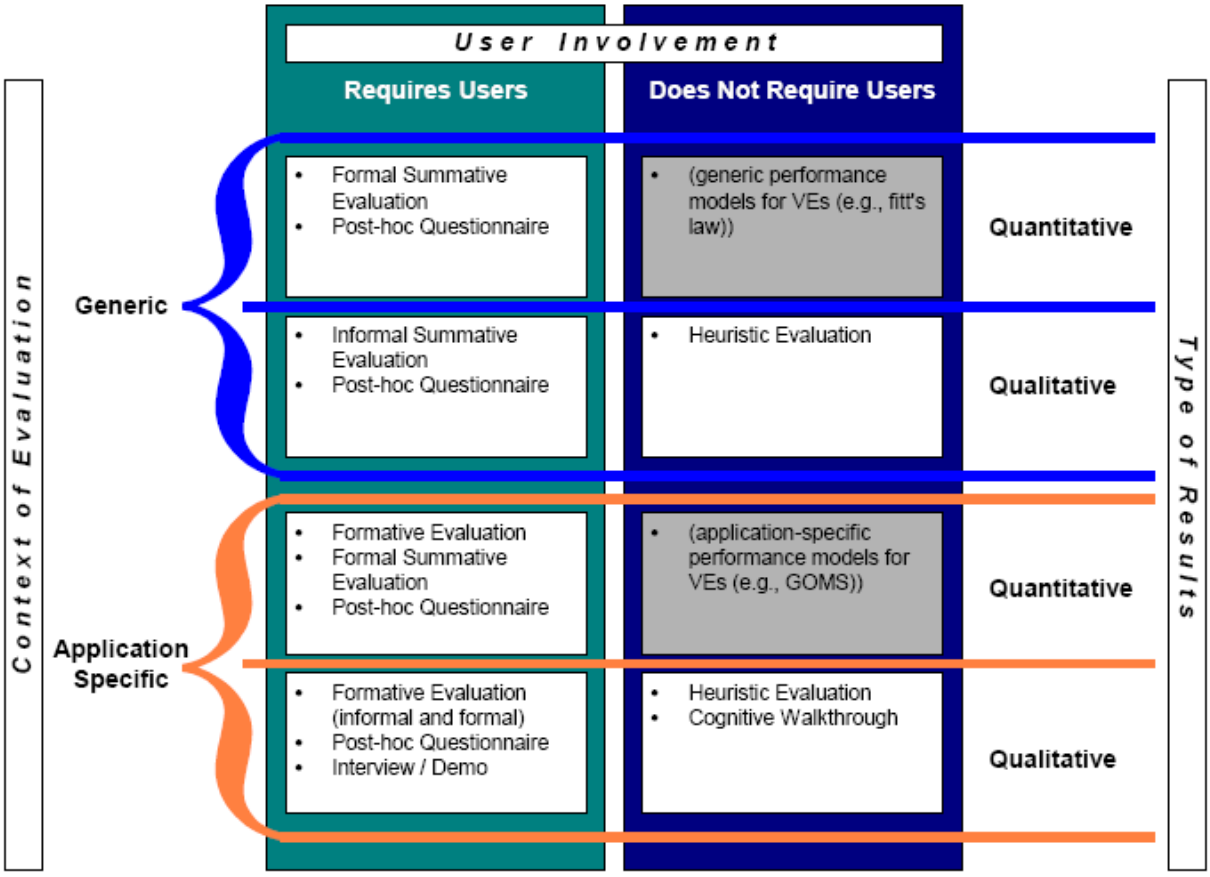

FIGURE 9 - A CLASSIFICATION OF USABILITY EVALUATION METHODS

Several other sources are cited in this report by bowman, which are also using both quantitative and qualitative data confirming the suspicion that they do work well together. The final choice and argumentation for this choice is explained in more detail in chapter seven.

Now all that remains is finding out when to apply certain evaluation methods. There is of course a plethora of research available on the process of designing useable interfaces. The usability design process is well known by now as is the fact that including users at early stages can greatly improve usability of said application. The iterative user-centered process is often used for these sort of things. Interesting to note however is that while this is the basis for VR evaluation methods as well, there are as of yet, two very distinct methods identified by Bowman. As mentioned earlier a problem in VR research is over-generalization of results, which is hard to do in VR because of the complexity and the unknown factors that still remain in VR. Because of this there are two very distinct processes being used to evaluate 3D designs. We will give a description of both processes. Even though our tests will be application specific, we are including the general testbad evaluation for later purposes, as we might get results that we want to test for their generelizability, instead of just assuming they can be generalized.

First of all we will discuss the test bed evaluation that is constructed for testing non-application-specific methods, of which there still are many today in VR being researched. By taking techniques outside of the context of applications and putting them in a generic context, and adding a framework for design and evaluation Bowmand and Hodges have hopefully created a method that provides us with systematic design and evaluation of techniques instead of only relying on experience and intuition. To have all this though, quite extensive testing has to be done. Since user issues and outside factors could play such a big role in VR research and these variables are sometimes hard to identify, the approach taken here is to vary outside factors as much as possible (down to the lighting model used for example).

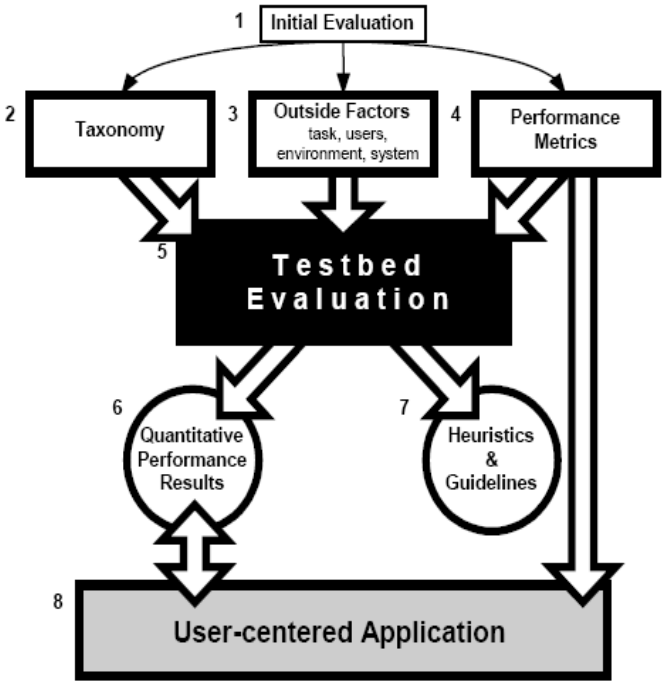

FIGURE 10 – A GENERALIZED TESTING FRAMEWORK: TESTBED EVALUATION

After the initial evaluation a taxonomy of interaction techniques for the task being tested is made, outside factors are quantified and metrics are defined. Using all of this information a testbed evaluation is done. These results can lead to techniques that can be used in generic user-centered applications and guidelines for the design of VR applications. The cost of doing this extensive testing is generally quite high so the benefit must be obvious before committing to this approach. However, the results may also be used as building blocks for complete applications, designed on the basis of the more traditional, sequational approach described next.

The application specific method mentioned in Bowman's book is actually researched by Gabbard, Hix and Swan in 1999. It is only quoted and explained in greater detail. We will give a short overview here of the method. However, as to what makes it exactly better than traditional methods might not be immediately clear. It deals with all the issues explained above and applies them to traditional usability engineering methods. To explain everything here though would require quite a lot of space and we will therefore suffice with this overview, since we already presented the issues.

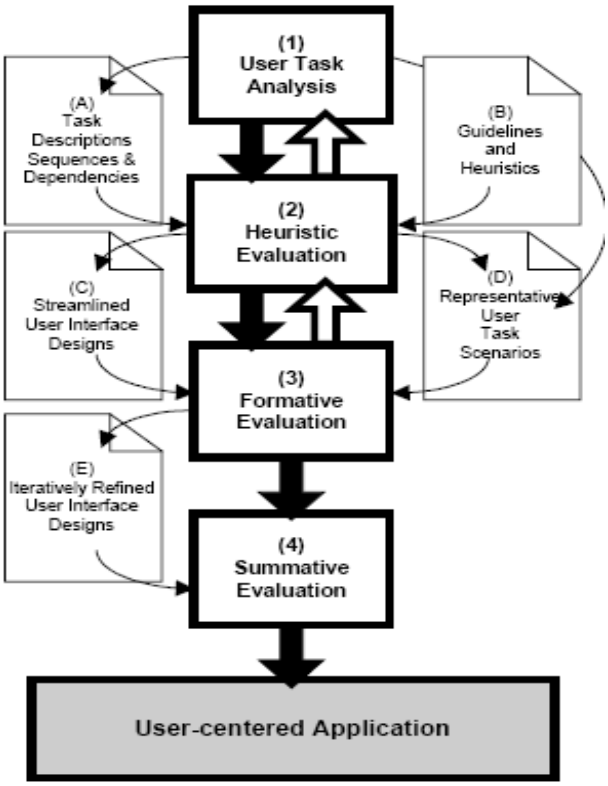

FIGURE 11 – APPLICATION SPECIFIC USABILITY DESIGN: A SEQUENTIAL APPROACH

As can be noticed there is the possibility of following just the traditional usability design path (indicated by the big arrows). However, by employing more techniques and specifying which ones, the large usability space for VR applications is quite well covered. All of these techniques have been around in the GUI field for years, the uniqueness herein lies in the breadth and depth offered by progressive use of these techniques.

We will use this sequential approach in the development of our own framework. However the earlier work done by Bowman concerning what tests one should run (observed in figure 9) is taken into account. The formative and

summative evaluation described in the sequential approach will be held using these tests described earlier, so both qualitative and quantitative tests in conjunction with one another.

Now that we have taken a look at some research about a big problem in VR, usability, which will play a big part in this project, we can take a look at how to actually design these interfaces for the public in the first place.

4.5 Design of 3D Interfaces

Designing 3D interfaces can be a difficult task. All the subjects discussed previously come together here. Specifying what it is you want to make, picking interaction techniques for this, putting it all together in an intuitive way and then testing it and making improvements where necessary will, if everything is done correctly, result in a usable VR application which serves its purpose. However, the road is long and a lot can go wrong, hence research on how to best design an application instead of just using our imagination and intuition is certainly called for.

4.5.1 Top-Down Design

Now there is not a standard way to formalize VR system design yet, but it is not without research. As mentioned previously, the process of designing 3D interfaces from a usability perspective has already been well researched. However the creation phase of the actual application can certainly use some fleshing out and it is here that there is a slight gap in research. There have been attempts to fill this gap and a particularly good one has been done by Parés & Parés (Parés & Parés, 2006). In their paper *Towards a Model for a Virtual Reality Experience* they describe this gap and their attempt to fill it. Parés & Parés describe this gap as "a deep theoretical gap in how Virtual Reality Experiences are modeled". What is important here is the mentioned *experience*. They define this experience as the perceiving of the Virtual Environment through certain interaction mechanisms.

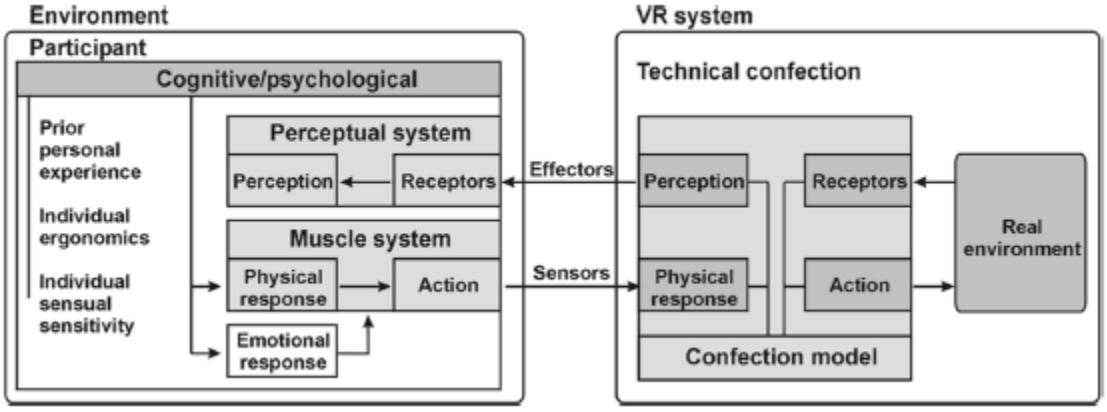

FIGURE 12 - HOW VR SYSTEMS ARE EXPERIENCED

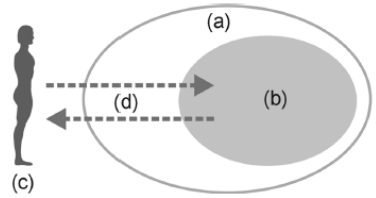

FIGURE 13 - A VR EXPERIENCE MODEL WITH (A) THE VR SYSTEM, (B) THE VE, (C) THE USER, (D) INTERACTIVE COMMUNICATION

As can be seen in figure 12 this involves quite a lot of fields. It is therefore not a trivial question and influenced by a lot of variables. Figure 13 shows a more simplified version for interactive VR systems. What is important to note here is the difference the authors have defined between a VR system and a VE, a Virtual Environment. A VE is defined as being the static environment that is modeled in the VR system. A VR system is than the evolving of the VE over time as, through certain interactions, regardless if

they are initiated by the user or the system, from one state to another. By manipulating this interaction between VE and the user (figure 13d) one can manipulate the experience a user has. A simple example is given where one would interpose a horizontal lattice work between the VE and the user, or even between the VR system and the user, and thus alter his or her experience (as opposed to a vertical lattice work). This is of course a very basic example; Parés & Parés follow up by showing more advanced interaction technique, one of which is very interesting and important for designing VR applications: mapping. If one for example maps the *field of view* of a device to be 180 degrees, the experience of a user would be very different than the more easy to grasp 90 degrees. However the user can also see more and have an experience that might simulate that of certain animals. One can also think of mapping of movement speeds. If one maps a virtual avatar to move very fast, the world in effect is experienced as a smaller world, even though it might be designed as a huge world. These are all important issues in trying to give a user a familiar experience.

As Parés & Parés are moving on and explaining more about how the experience is influenced by many different variables, they eventually arrive at a top-down method for designing VR applications. Even though it is still preliminary it does suggest a way of thinking where one first identifies main issues and makes sure that you know what the experience is that you want to deliver, before designing it. This way Parés & Parés hope to make the VR world less 'gadget-driven' and more 'user-driven'. A noble goal, but by now not very unique anymore. As may be learned from previous sections VR applications are slowly evolving from simply being research experiments to more matured applications, though much work remains to be done. In any case, their framework still provokes thought, and is also a nice way to keep track of all the issues one has to think about during the design of a VR system. The exact framework can be looked at in figure 14. It works on three levels, the application-level, the user-level, and the configuration-level. The reasoning is that the user experience should suit the application one is trying to make, and the configuration should make the user experience possible. Parés & Parés then go on to mention two very specific applications of this method, two theme-park attractions where user-experience is very important due to the nature of the application. One has however applied this model to every user individually while the other kept track of 'the big picture', which results in a better theme park ride (or so it is argued, this is of course a very subjective experience).

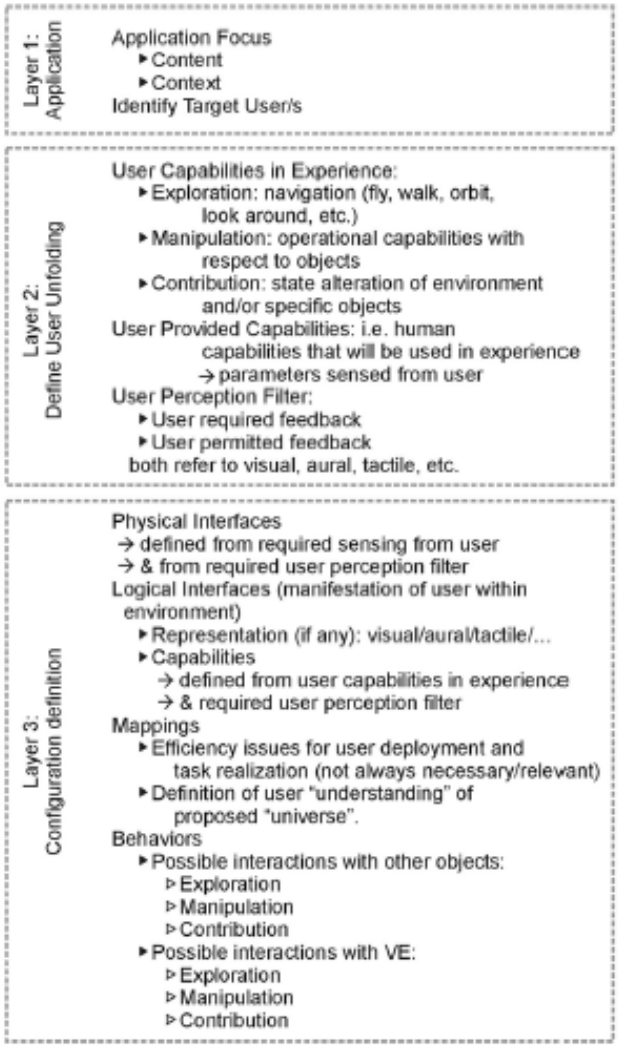
Layer 1: Application
Application Focus
▸Content
▸Context
Identify Target User/s
Layer 2: Define User Unfolding
User Capabilities in Experience:
▸Exploration: navigation (fly, walk, orbit, look around, etc.)
▸Manipulation: operational capabilities with respect to objects
▸Contribution: state alteration of environment and/or specific objects
User Provided Capabilities: i.e. human capabilities that will be used in experience
→ parameters sensed from user
User Perception Filter:
▸User required feedback
▸User permitted feedback
both refer to visual, aural, tactile, etc.
Layer 3: Configuration definition
Physical Interfaces
→ defined from required sensing from user
→ & from required user perception filter
Logical Interfaces (manifestation of user within environment)
▸Representation (if any): visual/aural/tactile/...
▸Capabilities
→ defined from user capabilities in experience
→ & required user perception filter
Mappings
▸Efficiency issues for user deployment and task realization (not always necessary/relevant)
▸Definition of user "understanding" of proposed "universe".
Behaviors
▸Possible interactions with other objects:
▹Exploration
▹Manipulation
▹Contribution
▸Possible interactions with VE:
▹Exploration
▹Manipulation
▹Contribution

FIGURE 14 - A FRAMEWORK FOR DESIGNING VR APPLICATIONS

This top-down design method is not the only design method being worked on. Bowman takes a more bottom-up approach, creating usable building-blocks first using the aforementioned testbed evaluation method, and by making these building blocks usable, hopefully making sure that an application built on these building blocks is also usable. These methods however are not mutually exclusive. There are certainly merits in taking a look at the big picture first, but being able to depend on well designed building blocks (perhaps in the future even standard solutions) to support the experience one is designing. We will therefore take a look at what Bowman has done on these building blocks and how one can design these smaller pieces of the application in a better way.

4.5.2 BOTTOM-UP BUILDING BLOCKS

Bowman has written several papers on this subject (Bowman, Koller, & Hodges, 1997)I & II, (Bowman & Hodges, 1997) (Bowman, Kruijf, LaViola, & Poupyrev, 2001), (Bowman & Wingrave, 2001) but most of them are bundled in his book (Bowman, Kruijff, LaViola, & Poupyrev, 2004). Since there are so many papers written on the subject (which are included here for completeness sake and for when the reader wishes to follow certain articles that he or she might find interesting) we will and give a general overview of the findings. An important paper however that one can use as a starting point is the one published in *presence* that gives a small introduction to the design of interfaces (Bowman, Kruijf, LaViola, & Poupyrev, 2001).

Subjects, or rather, building-blocks treated in the aforementioned papers are travel techniques, techniques for manipulation, menus (important for selection tasks) and viewpoint manipulation techniques (which could be seen as a subset of travel techniques). We will start by saying something about travel techniques.

Travel in Virtual Reality is defined as 'getting from one place in the VE, to another place'. The method in which this goal is achieved can however differ. One can think of many different techniques that can support the overall design of an application, especially in the case of a virtual museum. What immediately comes to mind is viewpoint manipulation, thus directly controlling the virtual camera so a first-person perspective of the 'virtual avatar' is achieved. Now there are many ways of doing this, some of them dependant on the interaction device, some more on the metaphor used. One can point to travel along a vector $\vec{p}$ where $\vec{p}$ is defined using something like a data glove, a headtracker (in essence using 'gaze-directed' steering), or any other tracker held in the user's hand. The vector $\vec{p}$ is then transformed into a world coordinate vector. There is also the method of 'camera in hand' where the hand *is* the camera through some sort of tracking. Furthermore there is the 'grabbing the screen' technique where users 'latch on' to a certain point in the screen in some way and move around relatively to this point. This has been found to be effective, probably because it is similar to our own method of looking around (our heads are also latched on to a certain point, and cannot make 360 degree rotations). In any case, as can be observed there are plenty of methods available for manipulating the viewpoint, and it looks like there will be more available as the amount of interaction devices increases. What is needed is a more general taxonomy of these methods so that a correct design-decision can be taken on the basis of criteria that are regarded as important. To address this problem Bowman has created a taxonomy which subdivides the problems faced with viewpoint manipulation in certain categories.

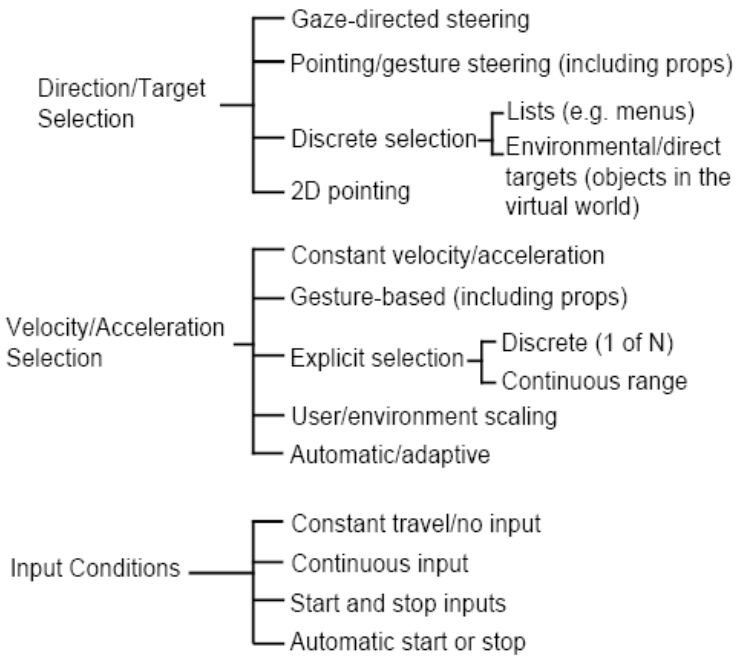

FIGURE 15 - BOWMAN'S TAXONOMY FOR TRAVEL TECHNIQUES

Now we can actually use several factors on which to classify our interface and can make a choice out of several alternatives. The ultimate choice will be dependent on more than just performance measured by bowman though since we will want to give the user a museum experience.

What might stand out is target selection. By selecting a target you are not actually using direct viewpoint manipulation, but rather selecting your target that you want to reach through some method, and then being transported there. Research by Bowman has shown that teleporting here, which might save time, is actually very confusing to users. However, a relative quick movement showing the path traveled to arrive at the destination was found to be more comfortable and allowed the user to quickly adapt to the new situation. This travel technique is called the *ZoomBack technique*.

Furthermore a museum could have a 'guided tour', where the user doesn't manipulate anything, but the viewpoint simply travels through the museum. An alternative to this to allow the user more freedom is presented by Bowman using something called *Semiautomated steering*. It attaches the user to a 'spring' which is attached to an anchor, which is travelling down a path. The user will then be pulled along but can still travel by himself (through direct viewpoint manipulation).

Now so far we have just presented travel techniques. However some of these conclusions also map nicely onto manipulation techniques, especially considering the mapping. Manipulation is often very basic and device dependant (and thus we will not say too much about it in the application design section) but the mapping is definitely part of the application design. A subtle flick of the wrist could rotate or move an object over a large distance. This however usually sacrifices accuracy according to research done by Bowman. The best way for people to manipulate objects is to map their hand movements 1:1 to a virtual hand, apparently also enhancing the user's sense of presence.

As for selection techniques, menus become more important. Unfortunately not much research has been done in this area yet. While Tromp et al. argue that since the user is in 3 Dimensional space, a 3 dimensional metaphor should be chosen to enhance this sense of presence, Bowman argues that for simple tasks, a traditional 2D selection menu often suffices, especially if the data to be browsed and selected is in a 'flat' structure (such as an array or a list). This seems to be good advice, as this is exactly something that can be experienced as 'contrived' by users. However, using new interaction devices, one can browse two dimensional data in a more intuitive way. Some suggestions with regards to gesture tracking and the data glove are made by Bowman.

Some interesting examples are the 1-DOF menu, which is simply representing a list in a three dimensional way. One can browse the list by rotating the hand, where the speed is dependent on the rotation. This gives a user a very direct amount of control without resorting to 'older' devices. It concludes that this is an efficient way, as long as the number of items is kept relatively small. Another proposed method using data gloves are TULIP (Three Up, Labels In Palm) menus, which have been found to be very effective. By projecting a menu option on the virtual hand at each finger, using the specific interaction made possible by a pinch glove (each finger can operate as a button) quick intuitive task-selection is possible. The subject then moves on to 3D widgets. This is only recommended when the menu has multiple dimensions which are better projected in this way, or if the interface somehow affords three dimensional movements. It does mention however it may give the user a more direct sense of control over the environment.

Now using all of these building blocks presented above, several design guidelines are also included in the book, which is very nice compared to the papers, since they are now gathered in one central place which makes it easier to use as a reference (One could say the book is more usable). A sort of obvious but important guideline, which works nicely with the top-down design method by Parés & Parés is the recommendation to match the travel technique to the application. This seems obvious, and it often is, but it is also important. More specific advice is given however. It is concluded that for goal-oriented travel, target-based travel techniques are preferred, while steering techniques should be used for exploration and search. Furthermore, if spatial awareness is important, one should use graceful transitional models, such as the ZoomBack technique. To help the user it is often a good idea to also provide multiple travel techniques, even if they serve the same purpose. It goes on to state that the travel technique cannot be chosen separately from the hardware used. This is of course true, if one considers target-based travel, a device that supports selection is needed.

As towards selection and manipulation tasks several design guidelines are given. Important is to not disturb the flow of interaction and prevent unnecessary change of focus of the attention. Tasks should not require mode-switching. Furthermore for menus, use an appropriate spatial reference frame, meaning that menus should be "in the right position". If a menu never gets noticed it will not get used.

Now having taken a closer look at all these important issues there were still some issues that did not fit in the previous categories. Since the virtual museum is supposed to also be a learning experience, next to being a "good" experience, we shall now move on to that subject, and propose an information structure that might support learning through exploring as well as some more research on the 'edutainment' experience.

4.6 OTHER ISSUES

This paragraph will deal with research that is only marginally related to our project. However the concepts explained here can be very useful in future evaluation of applications like this one or may be necessary to understand some of the proposed further research and to see what ideas we were playing around with when designing certain parts of the program.

4.6.1 SOMETHING ABOUT DATA STRUCTURES

To start we will introduce a certain structuring of data, which is as of yet not generally introduced to the public, but has some interesting ramifications, especially for visualization of data. The data-structure we are talking about is known as 'Zig-Zag' (or the more formal zzstructure) and is described by Ted Nelson (Nelson, 2004) in 2004 in the journal of digital information. It is especially interesting since it is argued that this design of data structuring (and any design in particular) cannot be argued about in terms of right or wrong, but this one is *good and usable* and the paper takes quite a postmodern approach on this.

The premise of the paper by Nelson (which in itself is an ironic name, article might perhaps be better) is that today's data structures and conventions are all based on some principal conventions, namely the 'simulation of hierarchy and the simulation of paper'. Examples of this are current day file systems that are arranged in hierarchies and literally as 'filing systems' where data must have a name to be represented. Relational databases are rectangular data, resembling paper once again. XML uses hierarchical structures and the output is usually something that is simulated to look like paper. Even programming is taken as example. The flow of programs or the grammatical makeup of a program is not apparent from the in-line view of code that is almost always presented when one is programming.

What is interesting for us is the mention of a 'view' on a data structure. We certainly do have a tendency to present data in a text-based, two dimensional manner as well as relations between certain data points. Spreadsheets, arrays (with table-view), online forms and other database-related applications seem to reinforce this point of view. What if we could view data in a completely different manner, where the data structure itself is something that can be explored and 'navigated' and where visualizations make certain connections between points of data obvious, without restricting ourselves to (for example) one dimensional menus. It is here that the article by Nelson ties in to our project by proposing a data structure on which very nicely visualized 'views' are possible. We will now give a short introduction of this data structure, and present the views mentioned in the paper.

The data structure is, for explanatory purposes, compared to a spreadsheet. A cell in a spreadsheet has at most 2 neighbours, in 2 dimensions. This system of cells and neighbours is retained in ZigZag, however more dimensions are possible. Also the 'ordered' notion where connections to other cells depend on its neighbour's connections in spreadsheets is dropped. These multiple dimensions are very abstract and hard to grasp. Where the spreadsheet had rows and columns Nelson refers to these structures in ZigZag as 'ranks' in dimensions. A simple example of a 'ZigZagged' spreadsheet is given in the figure below.

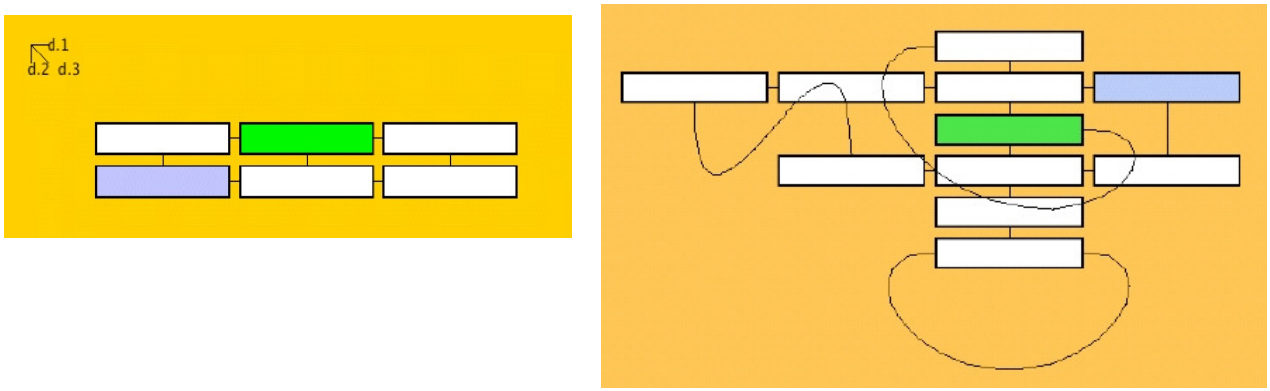

FIGURE 16 - A FIGURE SHOWING THE REGULAR SPREADSHEET STRUCTURE AND THE POSSIBILITIES WITH ZIGZAG

This structure is in essence based on three primitives, from which everything else is derived, and thus arriving at the desired simplicity of the system. The three primitives are zzcell, zzlink and zzdim, for the atomic cells, links between them and the dimensions in which it can travel. Now the system is kept simple by limiting the numbers of neighbours in each dimension to two, hence from a cell, one can only travel in two directions along a dimension, a positive or a negative direction. Links are untyped, symmetrical and there exist only one-to-one links. All the other common structures and relations such as many-to-many relations, trees, arrays, lists, etc. can now be represented by compositions of zzstructures.

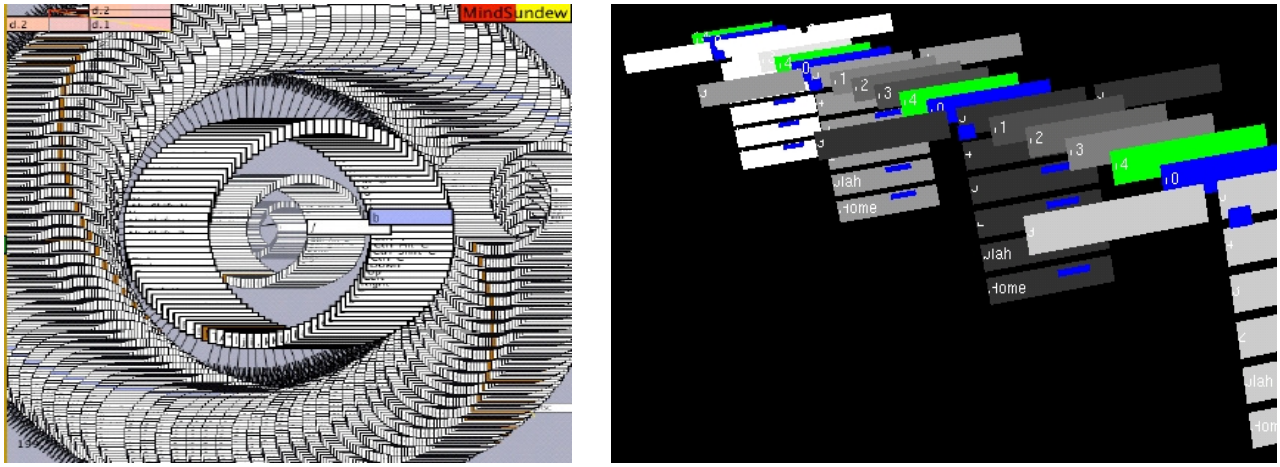

FIGURE 17 - TWO ZIGZAG VIEWS, THE MINDSUNDEW VIEW AND AN EXPERIMENTAL VIEW DONE BY NELSON'S TEAM IN OPENGL, MORE DETAILS CAN BE LEARNED FROM NELSON'S REPORT

Now as mentioned, this is quite abstract, and according to Nelson himself, much hands-on experience is needed to gain a full understanding of the possibilities and ramifications of this system. A simple example given is that, say we used a standard spreadsheet structure, we are irrevocably committed to putting data in 'quarters' and we might run out of space where to put certain data that does not fit the mold. The ability to 'just add a dimension' to your data here is what ZigZag represents.

What is of more interest to us are the views possible on this system, and how they translate to the 3-Dimensional world. Chapter ten of Nelson's article handles on how ZigZag could be used to visualize 3D graphics, or, how multidimensional Euclidean data can be visualized. By having a 3-Dimensonal point in zigzag associated in three dimensions to three real numbers, say *x, y* and *z,* you get the $\mathbb{R}^3$ space. Now by adding new dimensional data on your data point (or rather zzCell), say *p, q* and *r*, you would get $\mathbb{R}^6$ space. This can be extended to any dimension obtaining an $\mathbb{R}^n$ space. To visualize it one could follow a limited set of links (through a limited set of dimensions) and still have bits of 'data' correctly positioned, relational to each other. It is only not possible to visualize the entire structure at once without confusing most people since now each cell can exist in the visualization more than once (on multiple dimensions, for example both in the xyz view and in the pqr view). Solutions for this are still being looked at. We will show two screenshots of views currently implemented.

As can be observed in figure 17 the views can still be quite confusing. However exploring structures like this through multiple dimensions have a couple of advantages for our application. The relations between data can be apparent through just looking at the visualization, which might lend a new dimension to virtual museums by visualizing connections between objects. Furthermore, exploring along these 'multiple dimensions' can lead to unexpected results. Also, by being able to 'skip' through dimensions you might get to that one item that interests you, without having to go through an endless list of non-interesting items (compare it to jumping from any cell in

the spreadsheet to any other cell at once, if it is somehow related). There would be enough problems to solve such as how to remember from where you came, how to get to a certain point, but the idea itself is interesting enough.

However, since this concept is so abstract and complicated it will probably not be covered during this project, since it would probably require a project of its own. We will however keep matters such as how to navigate complex relational data in mind when designing the interface for the VR museum and that is the main purpose of having a closer look at this data structure as we just did.

4.6.2 The concept of Edutainment

A somewhat 'softer' concept than the previously discussed 'ZigZag' is that of edutainment, a topic that surpasses VR research and is broader than just that. Of course we would like to make our application both educational and entertaining, since that is generally why people go to musea, and we should try and surpass the normal museum experience in at least one of these aspects. First however we must try and get a feeling for what edutainment is, which we will do through some examples, and then on what sort of indicators there are that an application is in fact edutainment.

Extensive research on edutainment is done by Wiberg & Jegers (Wiberg, 2003), (Jegers & Wiberg, 2003) where multiple examples are given of *edutainment,* both websites and games, one particular game in an e-museum. In the introduction of her PhD thesis Wiberg stresses (as in previous papers) that for edutainment evaluation of the 'experience' rather than a functional evaluation is necessary. But what experience are we talking about here? To illustrate this we will now mention some example.

In (Jegers & Wiberg, 2003) a game is described that allows the user to manipulate a laser beam to get it to its end destination. These manipulations are done using 'real world' metaphors such as lenses and mirrors. Through this game a user playfully learns how light reacts to different stimuli and it is shown that users have a remarkable speed of gaining this hands-on knowledge. It is reasoned that this is likely due to the cognitive processes involved. A general statement on educational software can be found in this thesis claiming that "Good education software should be active, not passive, in that the learner should be doing something actively and not watching something passively."

Another example is the game 'Math Rescue 1'[8] where users can play a platform game. However to overcome certain obstacles the user will have to solve some basic calculation. A little story is included on how you are rescuing numbers that have been abducted and how the world cannot operate without these numbers. This game is obviously somewhat more aimed at kids.

An example where entertainment comes first and education second is the well-known civilization game series[9]. While originally meant as entertainment and trying to flesh out what it is that makes a game involving and fun, its designer Sid Meier also stated that he wanted to put some educational aspect into it, which is why every version of the game is shipped with a 'civilipedia'. In these games the player is tasked to create a flourishing modern civilization out of a tribe of nomads starting in 4000 B.C. It involves many real world buildings, world-wonders, technologies, religions, world-leaders and nations and contains extensive background information on each of these things and a small explanation of how they converted this to gameplay logic. This example shows a somewhat more passive educational aspect through the use of an enticing front-end.

[8] http://en.wikipedia.org/wiki/Math_Rescue as visited in July, 2008

[9] http://en.wikipedia.org/wiki/Civilization_%28series%29 as visited in July, 2008

As can be seen from these examples edutainment can take many forms. In some of these examples education is the prime aspect, in others it's entertainment. They are however not mutually exclusive and as such we should aim to try and teach people something, in an enticing way, such that the educational act in itself is also fun. A somewhat easier goal for a museum since one can assume that people going to a museum go there because they are interested in the subject and would like to learn something about it.

The entertainment experience however, needs to be defined some more and for this Wiberg presents research by Pine & Gilmore II which describes the experience realm. An area divided in four quadrants which cover the way we experience electronic entertainment.

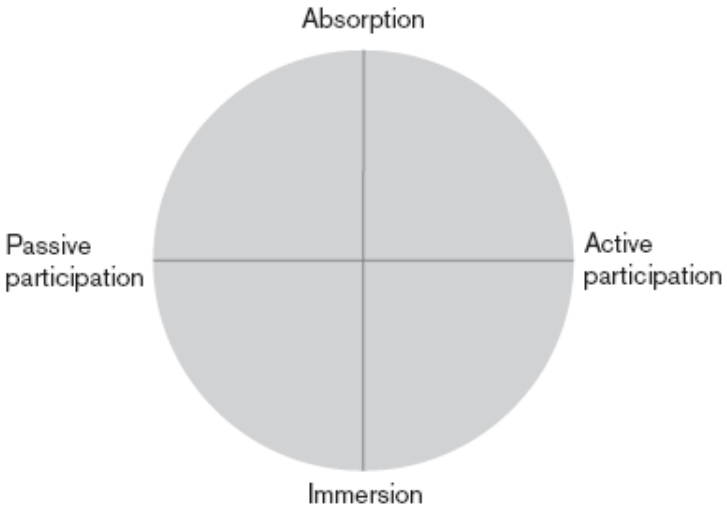

FIGURE 18 - THE EXPERIENCE REALM

Pine & Gilmore (according to Wiberg) have put entertainment and education in the top half of the experience realm, arguing that they both want you to absorb something, but once again, active participation enhances the learning process and should therefore be placed on the top-right, whereas the more passive absorption is seen as entertainment. Wiberg argues that entertainment can also require active participation and some games definitely seem to prove that point. To complete the 'circle' we will also mention that the lower half is covered by *escape* and *estheticism* (where the second is passive and the first is active).

An approach to defining fun is made, but ultimately abandoned since fun is in the end, a very subjective measure. However, based upon the heuristics review common for usability software, commonly using the ten usability guidelines by Nielsen, Wiberg has created a new set of heuristics referred to as the *funology heuristics*, we will now list them here.

1. Visual Impression vs. expectations: Do not let visuals make an impression that the interaction cannot meet
2. Exploratory design: Users should be enticed to explore the content
3. Playability-Gameplay: Visualize gameplay elements, otherwise the user will expect to be able to take actions he cannot
4. Durability and lifetime – amount of content: There should be enough content and not only for a small session

5. Coherence between design and chosen feeling: If you want to attain a certain feeling or mood, make sure the design of your application allows for this.
6. Clarity of genre – design for the right target group: Your information must be presented to the target group, this also implies finding out who your target group is.
7. Balance between information and entertainment: If entertainment is a goal, there should still be enough content to afford for this entertainment.
8. Originality and freshness: Information should be fresh and unique
9. Consistent navigation: An inconsistent navigation will influence enjoyment of any application by a lot.
10. General functional aspects: All the normal usability rules also apply here, as no unusable application is actually fun.

While these guidelines have been developed for websites, many of them are extendible to electronic applications. Using these guidelines and reviewing our design during the design phase, the suggested heuristics design in figure 11 for the sequential approach, will at least make us avoid some pitfalls that would make an application not fun. Hopefully that, combined with the active participation that is required in a web application, can interest users in the subject matter at hand, by having them make cognitive decisions about the content.

In the end we will choose to mimic an approach chosen by Wiberg to measure fun and try and define 'metrics' which together can make this Virtual Reality application fun, and thus breaking the subjective attribute 'fun' up in several smaller bits which might prove easier to measure. It will however remain hard to capture that elusive combination of factors which makes us think of something as fun and proving beyond any doubt that any application is 'fun' using only heuristic guidelines and expert review will very likely prove to be impossible; it is also highly unlikely that every test subject will unanimously agree that our application is fun.

4.7 Further Reading

The leading journal for Virtual Reality is published by MIT Press and is called "Presence: Teleoperators and Virtual Environments". This journal, which appears six times per year, contains many relevant peer-reviewed articles on various VR fields including Interaction design, usability evaluation and cultural heritage. Add to this articles found through more conventional means such as Google scholar[10], citeseer[11] and the university library and there is still an impressive array of articles available providing one with plenty of inspiration and ideas, as well as a source of ideas that have been tried and what works and also very importantly, doesn't work.

Next to these articles there are also plenty of books available. A somewhat heavy but comprehensive place to start is the "Handbook of Virtual Environments" which contains in-depth chapters written by a wide collection of researchers on almost every possible topic of Virtual Reality, from the way the human eyes perceives the virtual imagery to the evaluation of interaction techniques. It is edited by K. Stanney and has references to a very large collection of articles and books written by others.

Another good book if you are new to interaction design in VR is '3D user Interfaces' written by D. Bowman et al. It is more of an introductory book but still contains a plethora of research from others as well as practical tips and design guidelines that are used in the creation of VR interfaces. Bowman is a name that is more or less unavoidable

[10] http://scholar.google.com

[11] http://citeseer.ist.psu.edu/cs

when doing research in the Virtual Reality field and for an overview of the work done by him this book is also quite good.

4.8 SUMMARY

First of we showed some examples of previous work done concerning cultural heritage in Virtual Reality. We concluded that there is a relevant gap where our framework might provide some solutions. Furthermore the issue of digitizing information for cultural heritage was presented, with some examples and references to work already done on this.

We moved on to mention research done towards interaction devices and how there is not a clear ranking in these devices. We showed some advantages and disadvantages of some interaction devices. We were still left with a large number of known devices however; hence we have shortened the list somewhat and presented a list of devices we will consider for this project based upon several criteria explained there.

Since evaluation is a big part of this project we have shown some approaches to this for Virtual Reality, and have tried to explain the difference between traditional usability evaluation and issues specific to Virtual Reality. We have chosen for a sequential approach to evaluation, using a process presented in figure 11. We have extended this somewhat by using both quantitative and qualitative evaluation methods defined earlier in figure 10.

Using evaluation techniques and a design process we have described issues and techniques for designing VE's somewhat more. We described a top-down method focused on user experience which provided us with a framework in which to design our application. We compared this to the bottoms-up approach suggested by Bowman and concluded they were not mutually exclusive; hence we will use that research as well during design.

Finally, we have also presented some mildly related issues that might be important for future improvements on the framework and some research on the somewhat softer aspects of evaluating fun, where much more research still is needed.

In closing, all this research presented here will be taken into account during the design and creation of the Virtual Reality Application. However, since time and manpower is quite limited we will not be able to use all of the many elements presented in this chapter in the final design, especially considering it should provide an interaction framework and is meant as the basis for further work. Some research presented here could be called an 'exercise in academic thought' as we will not be able to reproduce everything that has been presented here. Those that are not included in the design however could certainly be avenues for further work, and they will be mentioned in chapter nine and are well recommended reading for anyone working in this area.

5 Design of the Virtual Museum Application

Having defined the problems we will face and the tools and related work that we may use to solve it we now move on to the design of the virtual museum framework. Our first step in this process, assuming that we follow a top-down design method, was to find out exactly what we would have to design and create, and which devices we were going to use to interact with our design. Based on that we have created some use-cases (or user task scenarios, if you will), which put the user-experience at first place. We have then designed an application that can support these tasks, and is easily moldable to fit the needs of the users. Note that we have followed the usability design process described in section 4.4.2 to arrive at a design as described in section 4.5.1 with three levels, the application level, the user level and the configuration level.

5.1 Initial User Study & Task Analysis

The problems we faced had to do with the choice of interaction device, the tasks that a virtual museum should support and how to evaluate if our design works as intended. Once we had fleshed out these tasks we could start work on a design to support these tasks, which we have later evaluated. The first step in our design process was than a task analysis. Furthermore at this stage in the process there was a large list of interface devices that were still considered feasible. This was narrowed down since we won't have the resources to implement solutions for every interface device nor the time to test them all.

To address these problems we have done an initial user study & task analysis, which gives us an idea of what users expect out of a device, and at the same time asking them about their expectations of a VR museum.

5.1.1 Method

The test was held at the Flexible Reality Lab in the Ingvar Kamprad Design Centre in Lund. For the interviews a quiet corner in the lab with a table and two chairs was reserved.

Before we would start the test with our subjects we found out something more about their backgrounds, hence a small knowledge test was done, as suggested by Lampton, Bliss and Morris (Kay M. Stanney, 2002). We asked the users about their experience using several advanced interface devices and their experience concerning 3D navigation. This was done since it might influence preference for a device later. Furthermore we made sure that the people we test are people who would be interested in going to a VR museum, hence they were not be technophobic, and already had an interest in visiting normal museums. To also gain expert opinions a computer scientist was invited to discuss technical possibilities, as well as a historian with large experience in museums and the transfer of knowledge to visitors. The eventual group consisted of three men and three women, aged 19 to 25, all sharing an interest in museums.

According to (Eberts, 1999) as quoted in (Kay M. Stanney, 2002) a user task analysis can be done using four different methods, namely:

I. Documentation review

II. Questionnaire survey

III. Interviewing

IV. Observation

Discussed there is that a documentation review is more suitable to find technical specifications based on legacy systems. This is true in some part, however one could also find inspiration in other applications; a documentation review is however, not in the scope of the initial user test and has already been covered in the previous chapter.

A questionnaire is generally used to evealuate existing interfaces. We have used a questionnaire about some features that might be added to the VR museum.

Interviewing was the main focus. During this interview we gained some more insight on what the subjects like about museums and more importantly, what are they missing. Some questions that we has the subjects answer were:

I. What sort of navigation is preferable, think about alternatives as 'strolling through a museum', very clear directives, teleportation to any (related) object, (related) objects teleported to you, etc.

II. Would it be important to be able to manipulate objects, or is it enough to control the viewpoint?

III. Would you mind 'educational' games or assignments during your tour through the virtual museum?

IV. Would a virtual simulation of an old event appeal to you, and if so, would you like to be able to interact with it, perhaps combined with an (educational) assignment (from question III)?

V. What sort of metadata would matter to you when you think about art objects, e.g. Creator, timeframe, location where it was created, style, references, etc.?

Of course it was considered important to also ask *why* users give certain answers.

We presented a list of features to users which we have asked them to rank on a one to five scale for appreciation and frequency where applicable. The list of features is as follows:

Feature	**Properties measured (on a 1 to 5 scale)**
Pop up information on a PDA	Appreciation, Frequency
Educational games	Appreciation, Frequency
Teleportation to objects	Appreciation
Teleportation of objects	Appreciation
Manipulation of objects	Appreciation
Directives	Appreciation, Frequency
Historical simulation	Appreciation, Frequency
Visible related data (e.g. from Wikipedia)	Appreciation, Frequency
Rating system	Appreciation

TABLE 2 - FEATURES PRESENTED TO TEST SUBJECTS

Using this data we were able to construct a prioritized list of task descriptions that are important to a VR application hence one of the goals of this initial user study was fulfilled. The task descriptions were then reviewed and used to construct user task scenarios, as suggested in the design process in section 4.4.2.

For gathering data on interface devices we have used a mock-up test. During this mock up test we have asked users to use two of the devices presented in section 4.3.3 and show us how to use them in a natural way. We mimicked three tasks for this: navigation, selection and manipulation. Special EON simulations were selected that to mimic these tasks. For navigation we used a small exhibit area which the user was asked to 'navigate'. Selection was tested using the presented exhibits. For manipulations a modeled sword was used. The choice here is based on an object which is supposed to be 'waved around'. One note however, after the first test the PDA proved so cumbersome for navigation that we did not test it for this task anymore; hence we will leave it out of the equation later.

Since this was still the initial phase nothing actually worked during the mock-up test. However the person responsible for the test made sure the simulation did respond to the actions of the user, to see if the effect in the world is really what a user was expecting it to be. This required careful observation of the test subjects as well as establishing a good rapport with them to gather if the actions taken are indeed how they meant them to be. Care was taken in noting the response-time of subjects. Taking a long time to answer how 'natural' a certain action with a device is, might indicate that it is not so natural at all.

The test was done in a 4-sided CAVE, providing the users with a stimulating experience and providing them with a serious high-tech atmosphere to make sure they take the test seriously. Also movement through 3-Dimensional space might be better sensed in such a highly immersive setup.

Once the mock-up test was done users were asked on how they would rank the devices they used. This allowed us to see how they compare against each other and if there are devices users always seem to prefer, regardless of the alternative. By restricting the amount of devices to two for each user they did not mix experiences or forgot about the first device they used if they were to use a lot more after that one. Occasionally however if it seemed interesting a third device would be added (for example when an interesting point to compare two devices on would come up).

5.1.2 Result

We will now present the results from this initial user test. First we will give a short summary of the answers given to the questions. We will not present every user's viewpoint as we would quickly run out of space, but rather the common denominator between the subjects. We will then present the questionnaire results, the average vote and the standard deviation. The standard deviation is included to get a feeling for how much user's views differentiated on these points, a high standard deviation indicating that users were split on this issue. After that we will give an overview of the devices, issues that were encountered, and how they compared against the other devices.

- **What sort of navigation is preferable, think about alternatives as 'strolling through a museum', very clear directives, teleportation to any (related) object, (related) objects teleported to you, etc.?**

The first test revealed a strong liking for the fact that you can 'stroll' through a museum and end up at exhibits that you might have missed in a targeted search. This question was repeated later to each user and they all confirmed that this is one of the things they like about a museum. However a lot of people indicated interest in the concept of 'teleporting'. While it is shown by Bowman that teleporting can lead to a decreased sense of direction and orientation (Bowman, Koller, & Hodges, 1997) most users felt confident they would soon know where they were as

all it would take was a simple look around the room. However when asked if they thought that this would also be the case if they teleported forward (to a non visited part of the virtual museum) instead of backward they were not so sure. It seems that being able to quickly go back to an already visited part is appreciated though, since it saves you a walk. The teleporting of object was often a very abstract concept for most people and depended heavily on the implementation. While some users were mildly enthusiastic about being able to teleport some objects together to be able to compare them none reacted with very strong enthusiasm.

Important here is to note the amount of freedom people wanted. Generally people would report that being able to 'fly' would be interesting, and especially manipulating the viewpoint in 'unusual positions' was seen as an interesting addition, for example if someone wanted to get a view on an exhibit from high up. This coupled with the historical simulation (for example, getting a bird's view on a medieval square) generated quite a lot of enthusiasm.

As to directives, most people agreed that they would need to be incorporated even in the 'virtual world' and not just on your 'PDA' or on demand by pushing a button. None argued that directions in a virtual museum would be experienced as annoying. The presence of them would even be appreciated, as long as it is kept to doors between rooms. What users would like to see is a sort of roadmap, just as in the IKEA, which shows what rooms are still to come, instead of just the next one.

- **Would it be important to be able to manipulate objects, or is it enough to control the viewpoint?**

There was only one user that was happy with just being able to only manipulate the viewpoint. The most heard argument was that they actually expected to be able to do this in a virtual museum, since it seems like a very basic thing which normally can't be done in a museum. Also rotating an object took less time than walking around it or moving above it and looking down on it. However, it was also noted that this was not a make or break issue.

- **Would you mind 'educational' games or assignments during your tour through the virtual museum?**

This issue was always appreciated, though in different measures. While everyone thought it would be a fun way to test your knowledge, not everyone wanted games in the same amount or in the same way, moderation was always good, they'd rather have too little than too many. A quiz was already enough for most persons, but they would like to be able to review their scores at the end and perhaps compare them with their friends, which suggests personalization and a coupling with some sort of database. Complete games as produced in the entertainment industry were to be avoided, since that is not the aim of a museum. Most users felt that if challenged, they would also retain information learned during their visit better. The concept of getting questions and having to play games during a museum visit did have a considerable impact on the presentation of a museum according to one user. It would definitely make it a less monotonous experience. Only one user didn't care for the games and said that he might even take a museum with games less seriously.

- **Would a virtual simulation of an old event appeal to you, and if so, would you like to be able to interact with it, perhaps combined with an (educational) assignment (from question III)?**

This one feature was really appreciated since this is something that you cannot do in every museum. There are some open air museums dedicated to these scenes but most are either small or have limited access. A scene that could be navigated freely while exhibiting certain objects was something that appealed to everyone unanimously. Another point where everyone was also in agreement is that they would rather have one big simulation instead of a couple of small ones. Later on it came up that this might actually be an interesting way to exhibit certain objects,

for example showing old tools in their natural environment. By seeing objects in their context it was argued that one might learn more.

The interactivity of such simulations was another matter. People were split on this issue. Two persons thought it might become a bit much, especially if it wasn't exactly clear what you were supposed to do. One of them feared it might become too much like a computer game. Others would like it if there was some interactive element like opening doors, small assignments (like "find the smith's hammer") or the like which gave them more of an 'objective' than just random exploring.

A suggestion made by one person was to make the scene 'alive'. Have animated elements and in the best case, actual people walking around in the scene. He realized that virtual reality probably wasn't on that level yet but some semblance of a lively scene (moving objects, sounds) he would certainly appreciated. When others were confronted with this they were quick to agree that it would be a good thing and it would enhance their enjoyment and sense of wonder.

- **What sort of metadata would matter to you when you think about art objects, e.g. Creator, timeframe, location where it was created, style, references, etc.?**

Every single user indicated that they were interested in the background of an exhibit. What they meant by that is the reason for its existence. Why it is shaped as it is and what is special about it. This is a bit vague but the most important features were the timeframe and the place when asked to be somewhat more specific. The reasons were that these features, according to those asked, often have an important bearing on the reason that something was made.

It is interesting to note that the creator of an exhibit was deemed less important. The first user that was questioned actually said that he usually didn't care if it wasn't a big name. This was later confirmed by everyone, though most wanted to be able to know since that is a way an artist could become a 'big name'.

As to how to find out more about an object, there seemed to be some interest to being able to browse through hypermedia structures, however, it should not be like Wikipedia since than one might as well use a computer at home. Two or three levels of depth is more than enough for most people (e.g. if you could find out something more about any of the related features).

- **Do you ever miss anything when you visit a museum?**

The first answer to this question was that the presentation of a museum is often quite monotonous and somewhat boring. You just look at exhibits and read something about it. Some more variety was wished for as I elaborated on this point to the other users. A proposal was the interactive games, short movies and audio, not necessarily at the same time but mixed, to have a varied experience. It was expected that this would most definitely enhance the experience of a visit.

A point that came up a couple of times was that museums sometimes lacked places to sit. A lot of walking was done during a museum visit but sometimes people would just like to be able to sit down and 'look at a room'. This might translate to a VR museum by making the interaction operable in a sitting position.

A good point was made concerning the size of the collection. One user mentioned that sometimes he was really interested in a certain artist, timeframe or something of the like, but the museum's collection was limited. He would like to see bigger collections since in VR scaling costs and space are less of a problem. Another user

elaborated on this point by stating that she didn't necessarily want a bigger collection but a more personalized one.

All this is quite a lot of information, from which we will draw some more summarized conclusions in chapter nine; we will now provide an overview of the ratings given to certain features of this museum:

Feature	Appreciation		Frequency (if applicable)	
	Average	Standard Deviation	Average	Standard Deviation
PDA Popup	2,4	1,29	2,5	1,29
Educational games	3,9	0,9	2,9	1,34
Teleportation to objects	2,8	1,3		
Teleportation of objects	2,6	1,14		
Manipulation of objects	3,8	0,84		
Directives	4,2	0,76	2	0
Historical Simulation	4,6	0,55	1,6	0,48
Browsing related data	4,2	1,04		
Rating system	2,8	1,64		

TABLE 3: RESULTS OF THE QUESTIONNAIRE

Conclusions can be found in section 5.1.3

For the devices, the following results were observed

WiiMote

The WiiMote was generally liked. It was tested five times, since initial interest in the WiiMote warranted comparisons to more devices. Buttons were often used instead of gestures to navigate (three out of five times). The times gestures were used users occurred unsure about which gesture to use. For selection opinions differed. Almost everyone (four out of five) used the A-Button, however to indicate what was selected it was split on the issue of a 'crosshair' mode vs. a cursor. For manipulation the expectation was always that the manipulated object reflected one on one the orientation at which the WiiMote was held.

Stereo Camera/Motion Tracking

To navigate using motion tracking was quite a surprise for most people. It took a long while to come up with gestures to manipulate the camera. Pointing as described in (Mine, 1995) was never thought of and mostly it would involve cumbersome gestures such as 'flying as superman' in a direction. There was also a problem which

was described in (Kjeldskov, 2001) where a user would turn towards the end of the 4-sided cave and then wouldn't know how to continue turning. Walking in place worked a bit better with some mixed responses, hesitantly positive though. Turning remained a problem. Selection was almost always done using pointing, or grabbing in the direction of something. Manipulation would then be done using the hands, as if one was really grabbing the object and moving it. Navigating a menu showed more promise, using gestures to go left, right, up or down in the menu.

Data Glove

The data glove suffered mostly the same problems as motion tracking for navigation. The difference between a data glove and normal motion tracking did not seem apparent to most people. When asked for the reason why they were uncomfortable the response was that instead of using vague gestures, a control-box such as the WiiMote or SpaceBall made them feel more in control. For selection, results were once again similar, pointing and grabbing were very popular. Manipulation was often done rotating the hand while grabbing. When confronted with the fact that objects might not always be shaped in something that can be grabbed users would still prefer grabbing.

SpaceBall

Results for the SpaceBall were quite promising. Everyone was very quick in imagining on how to use it. It was used quite successfully for both navigation and selecting; not at the same time however. Depending on the amount of DOF controlled persons would either move the SpaceBall forward or roll it; research was inconclusive however on what would be best, since people found limitations in the Degrees of Freedom between two and six too hard to imagine. The SpaceBall does apparently afford squeezing for selection instead of using a button; this is once again the 'grabbing' gesture repeated. However after finding out that the ball was not squeezable both users who tried squeezing first quickly switched back to the buttons. Manipulation was also mapped one on one to the object. Because of the promising results with both the WiiMote and the SpaceBall they were compared twice, with two different results, hence it is inconclusive as of yet.

5.1.3 Conclusions

5.1.3.1 Interaction Devices

In the final application we have not implemented a 'walk-mode' that utilizes a stereo camera or the data glove. Using gestures to walk was unnatural to users; especially for changing direction without actually moving your body. Instead efforts were focused on doing this using either the SpaceBall or the WiiMote. However, gesture recognition showed some promise for navigating as well as making people enthusiastic, and as such it was made possible to browse content or rotate objects using gestures in the final application.

A WiiMote showed tremendous promise, though not because of the reason initially suspected. The expectation was for it to be appreciated because of the familiarity, however, the functionality and the ergonomic design were the most given arguments. Using a WiiMote in application is going to need more testing than just this framework and as such it was also a goal to create an application where the WiiMote is linked in an easy way, so that others may experiment with this as well.

The SpaceBall is an intuitive device that affords moving it, and also squeezing. It might be considered in the future to make a 'squeezable' SpaceBall so users can 'grab' things, which might work really well to switch from navigation mode to manipulation mode in a natural way. In any way, it is a comfortable device to be used sitting, but not so much to be used standing. In the actual usability testing it remains to be seen if users prefer standing or sitting.

5.1.3.2 Relevant User Tasks

As can be observed in table 3, the historical simulation was by far the most popular feature with almost everyone agreeing on it. As mentioned one big one is seen as better than many small ones hence when designing a virtual museum, effort should go into creating a large simulation of a historic environment, preferably alive. One of our user-task-scenarios is modeled after this. The historical simulation is probably not as alive as people would like it to be but this is definitely a recommendation when this framework is finished.

Seemingly not very high tech things such as browsing related data and getting clear directions by just looking around in the museum were also highly rated. This is interesting to see as sometimes high-tech gadgets can actually draw people someplace but as can be seen relatively simple things are greatly appreciated. Furthermore the interactive elements such as manipulation and games are seen as fun and entertaining. During the discussion it was often commented on that collaboration on games would be the best. An intuitive and non intrusive way to play these games was therefore introduced into the application. Since manipulation is a well researched issue which was not make or break, and also ranked below browsing related data, playing a game, navigating and the historical simulation we did not include this in the framework.

Navigating around the environment was seen as a very important task, especially the manual viewpoint manipulation. People do want to be able to 'fly' and look around, hence more than 2 DOF is called for in the viewpoint manipulation. Rolling was left out.

Finally, the rating system which would allow us a better personalization of the data was quite widely discussed. Some people thought it was really good, but they seemed to like the possibilities that this rating system helped create, rather than the rating system itself. Some more research would be needed on this subject, especially passively gathering preference indicators which can be used to build a user profile, as discussed in section 4.2.

5.2 User Task Scenarios

5.2.1 Considered Tasks

For the user task scenarios used to base the design upon several tasks were considered based upon their rating in the initial task analysis. There was however not enough time available to implement or test all of these high ranking features, especially considering we are constructing an interaction framework and not a full-fledged application yet, hence some had to be dropped. This was mostly done on the basis of them being not a 'make or break' issue or not being interesting enough for research. We will briefly describe them and the reasoning for including them or dropping them:

Taking a stroll through the museum

Based upon the user task analysis previously performed this seemed a paramount function to support. Users appreciate the freedom of navigating around a museum themselves to explore and look around, in a manner of their choosing. While they appreciate a certain structuring in the exhibits the exploration part is very important to their experience. This of course means we have designed a user task scenario for strolling.

Setting up a custom ordering/Choosing from pre-made orderings

This is a task that was dropped. Users did appreciate having an ordering in the content somehow; however they did not explicitly indicate the wish to be able to customize it themselves. A compromise might be found by having them choose from a pre-selected set of orderings, retaining the expertise a 'virtual curator has' while offering a choice, and this could still be done once this framework is extended. However, to seriously test this one would

need at least two distinctively different ordered collections, to show users how their choice influences the experience, hence have not implemented this option.

Exploring content related to an exhibit

As indicated this is an important function for any museum. Information that is somehow related to an exhibit should be fully explorable and preferably relations should be made visible. Important here is that different forms of content are available to explore to make the presentation less 'dry'. This task was also ranked very high and thus included in the design of user task scenarios, considering it is also a novel task in virtual museums.

Playing a small game

Users all seemed to like this suggestions hence a small game is implemented. Another argument to support spending time on this is the fact that it might make information retention better and thus serve the educational purpose of a museum that much better. It should not be too complicated though, as we are not creating a game but a museum.

Manipulating objects

As previously mentioned this task was not included. Manipulation was ranked the lowest of these considered tasks, though still considerably high. People however did mention they probably wouldn't miss it if it wasn't there, but would enjoy it if it was. Another reason is the amazing amount of research done on manipulating objects. While it is interesting to see if it really adds something to the application in this context there is nothing specific about manipulating objects in a museum hence if one ever wishes to add this in the future, it can simply be implemented. The only thing that would need further research is how the interaction device would support it, which is already available for the SpaceBall (included in example applications with SpaceBall software and popular CAD applications).

Exploring a historical simulation

Considering the popularity of this task we had to implement this, even though it could be a lot of work modeling. The impact it can have on the experience of visiting a museum is remarkable and furthermore we could even consider using a realistic historical simulation as exhibition area. However if we wish to ask this to users they would have to understand what is meant here, hence this was included.

5.2.2 Setup of the Scenarios

We will describe the scenarios in several terms. Since we are describing the experience a functional requirement will not be enough. To overcome that we will add 'important issues' to each scenario which contain implementation suggestions on how to achieve a certain experience. Furthermore we will make a list of tasks that a user should be able to do, and what the system's responsibilities for those tasks are. The administration of a user task vs. system responsibility has been taken from the book "Exploring Interface Design" (Silver, 2005), and seemed like a good idea since we will need this information once we start the implementation.

5.2.3 Scenario One: Taking a stroll around the museum

User Tasks

1. The user is able to move his viewpoint (the camera) from one point to another
2. The user is also able to change the viewpoints orientation by pitch and yaw, roll is deemed unimportant.

3. Preferably the user will be able to do (1) and (2) at the same time to approximate the feeling of walking whilst looking in another direction than the direction one is walking in.

System Responsibilities

1. The system allows for camera locomotion through some interface device.
2. The system supports different orientations of the camera through some interface device.
3. The system can keep track of user's position and orientation.

Important Issues

1. The user must have a sense of freedom; he must not feel restricted or as if on rails with 'invisible' walls so the user gets a sense of being able to explore the museum, just like a real museum.
2. It must be *easy* to look around while walking for the user, after some practice this should be an almost unconscious action.
3. It should also be easy to orientate the viewpoint, resulting in interesting views on exhibits.

5.2.4 Scenario Two: Exploring content related to an exhibit

User Tasks

1. The user can select an exhibit.
2. The user can browse through objects related to an exhibit through some navigational structure.
3. The user is able to select an object after navigating there to access options related to this object and if deemed necessary browse content related to this selected object again.

System Responsibilities

1. The system can generate a menu-structure based on data that is related to a selected exhibit.
2. The system is able to place this menu in the 3-dimensional space where it is visible and accessible to the user.

Important Issues

1. The selection mechanism should be understandable without much conscious thought.
2. It should always be possible to return to the previous item one was browsing.
3. The selectable object should be placed within arm's reach of the user if possible, to afford the 'grabbing' or 'swinging' gesture that might be implemented later, and to be visible.
4. The user should not be presented with an insane amount of options or related objects to choose from, the 7 ± 2 rule seems like a good guideline here.

5.2.5 Scenario Three: Playing a small game

User Tasks

1. A user can access some sort of game, which is related to the exhibit, probably through said 'related content' menu system.
2. The user can play this game and receive feedback whether or not his try at the game was successful.

System Responsibilities

1. The system should keep track of said user's success rate.
2. Preferably the system should be able to store 'scores' between sessions to build a personal profile for users.

Important Issues

1. Results should be registered and presented at the end to provide a user with incentive to do well.
2. These should not be in the tour too often; a couple of times is okay.
3. The quiz-questions should be well designed, to challenge the user to get to know more about what he's looking at.
4. If possible, scores should be comparable to 'friends' or averages to create a sense of competition.

5.2.6 Scenario Four: Immersing yourself in the 'Historical Simulation'

User Tasks

1. The user is able to somehow enter a historical simulation; this can be from the original virtual museum environment or the original environment already.
2. In this simulation normal navigation is still possible to attain interesting perspectives.

System Responsibilities

1. If two different simulations are used here seamless switching between the two should somehow be possible
2. The system must have the capabilities to show large, complex environments.

Important Issues

1. There should rather be one large simulation than several small ones.
2. A simulation can be livened up a lot by the use of sound, this is relatively simple to implement so it should contain a few sounds.
3. Things in the simulation should be happening without the user having to start everything; this will make the simulation feel more 'alive'.

5.3 The Design of the Virtual Museum Framework

After describing these scenarios and getting a sense of what it is users would like out of a virtual museum by actually asking them we could start to create a design which addresses these problems. In summary, we must make a design which makes it easy to 'walk around' some environment; this environment should be what users expect out of a museum. In this environment users must be able to select exhibits somehow and have a menu structure which can be used to browse through related content. Preferably the same structure (for the sake of consistency) is used to access certain games or tasks associated with these items, unless this task warrants its own specific interaction style, and not a menu (such as manipulation). Using these tools the users must be able to seamlessly go into a historical simulation of some historical scene which is relevant to the museum.

5.3.1 Environment

Users have indicated that they would like a showcase of exhibits that is designed as in a normal museum. This suggests a normal exhibition room might be where people feel most at home. A normal museum is usually a large

building or space subdivided into different areas/rooms, where routes through these areas/rooms make you follow a certain theme. This was appreciated by most people so for this framework we have created an exhibition room with a couple of exhibits; the theme is neutral which might be an advantage when adding content of a different theme. Later on when this framework might be used to create larger exhibitions some more experiments might have to be done to get more details on the size and amount of rooms people are willing to go through, we have tried to include preliminary numbers though.

There also is a 'historical simulation'. This environment was not designed ourselves; instead have reused an old master's project that was created at the VR Lab in Lund. This is an outdoor environment which is quite large, with a church, a graveyard and surrounding cottages.

5.3.2 Overall look and feel

To support the feeling of being in a museum we decided to base our museum on objects found in a real museum with a real background instead of a collection of test objects. For this we collaborated with the Kulturen museum in Lund. Models and information about those models is based on objects found in this museum which we were allowed to photograph for this project. Even though we are just creating a framework, providing this look and feel was considered important in the upcoming usability tests, since we wanted users to feel as if they are using something that can be a real virtual museum someday.

5.3.3 Navigation

The design for the navigation should focus on allowing users to stroll around the museum easily, allow them to reach objects quickly and effectively and on getting an interesting viewpoint through manipulating the orientation. To implement this mode of navigation we used two techniques: ZoomBack and Manipulation of the viewpoint. Guided tours, tours on rails or spring based boat rides, as discussed in chapter four was not used since they constrain the user too much to a certain course. However if this application is ever extended with a guided tour it might still be a good idea as an addition. These base elements (ZoomBack and Viewpoint Manipulation) should not be taken out though.

5.3.3.1 ZoomBack

As already mentioned, navigation should be done primarily through manipulating the viewpoint. However, since users also have the ability to select objects and have indicated that teleporting would be interesting we can insert something extra here. As mentioned in chapter four there exists a technique that 'teleports' users but does show the animation of reaching the next point called the ZoomBack technique, described by Bowman (Bowman, Kruijff, LaViola, & Poupyrev, 2004). Hence if we define a viewpoint on every exhibit with both position and orientation (in terms of the absolute world), let's call this P_e , with position *x, y* and *z* and orientation *h, p* and *r* (for heading, pitch and roll) and we call the position of the camera P_c with the same properties, there should be an animated transition possible from P_c to P_e , hence the user is rotated and moved to an absolute position in the world which is pre-defined by (for example) a virtual curator. This movement is observed on screen as the viewpoint is moving through space. The duration of this animation should not be too long as this will slow down interaction, and research has shown that even if the animation is 'very fast' (the example named is ten times as fast as regular movement) a user will still have a good spatial reference. Since we are not including 'roll' in the simulation it is advised to keep the value for *r* at 0.

Using this technique allows persons to get to exhibits faster, as well as acting as feedback when one selects an object. This prevents users from getting bored if they just want to reach an object without moving the viewpoint there themselves, while still retaining the possibility to do so. It also allows for less accurate navigation to reach exhibits, as one can simply select them to be taken there.

5.3.3.2 MANIPULATING THE VIEWPOINT

Manipulating the viewpoint means that the user is able to change the *x, y* and *z* parameters for position as well as the *h, p* and *r* parameters. In effect this means that next to traditional 2DOF navigation, as for example described in (Wallergård, 2007), users will also be able to go up and down, and look downward and upward. To keep navigation consistent we would advise that a user is always moving in the direction that he is facing. That means that if he is facing downwards, he will also be moving downwards at that angle (e.g. if his pitch is -45° he will be moving in that direction, hence affecting both his (x,y) position and his z position, as opposed to just x and y in traditional 2DOF movement). In effect this means that the user is always moving and yawing within a 'plane'. The orientation of this plane is the same as the orientation of the viewpoint.

A more detailed explanation can be observed in figure 19. The red plane is the plane in which the user is moving. The grey line shows the orientation of the user (the direction where he is looking), thus the 'relative' y-axis. The red dot shows the position of the camera. If the user now moves forward he will move along the light grey line. If he rotates he will 'rotate' on the plane and might up in a position that relative to the absolute coordinate system is rolled.

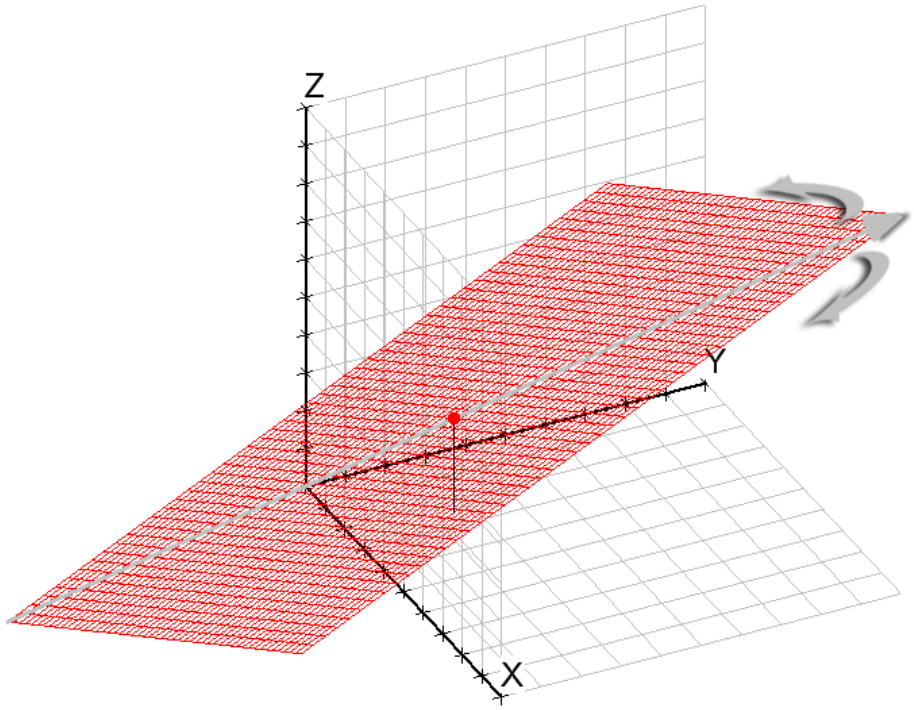

FIGURE 19 - THE NAVIGATION PLANE (THE RED PLANE) IN THE ABSOLUTE WORLD COÖRDINATE SYSTEM

Now by making sure the system retains control over this 'reference' plane, which we will refer to as relative coordinate system from now on, one can control in which plane the user moves and also how he can orient himself. Using this control it will be easy to later limit the Degrees of Freedom a user can manipulate to the

application's need later, if user testing might proves this necessary. Another advantage of this is that placing objects relative to the user becomes easier, since it is just a matter of putting them in the right place in the relative coordinate system. All that is needed is to continuously update the orientation and origin of the relative coordinate system as the user is moving to correspond with the camera position and orientation. For example if one would no longer update the pitch or roll of this relative coordinate system the plane would always be horizontally and the user could no longer move diagonally up or down, but the user would still retain the ability to look up or down. Note here how in the previously mentioned example the camera orientation is *not* the same as the relative coordinate system orientation since we 'unlinked' them to obtain that specific result.

This suggestion is already mentioned in the design phase since the way this model works should always be implemented, regardless of the chosen implementation, because there is not much research done yet on manipulating something else than 2 DOF or 6 DOF. Since we already concluded that roll was not important we are left with 5 DOF, and this might still change depending on the user study. There are many uncertainties and using this dynamic model of viewpoint manipulation will allow us to make changes later.

5.3.4 Selection

For selection tasks (selecting objects, selecting related content) we have designed both a manner of selection in the program, and a menu structure that supports selecting content and at the same time is able to show complex relations which might be added later. We will start by designing the mode of selection in the program.

5.3.4.1 Object Selection

A natural manner of selection for persons was to 'select what they were looking at' (by executing the appropriate 'select action') which in our navigation model terms is translated to that which is in front of the camera, thus, the first object to be intersected by a vector *v* that 'shoots out' along the y-axis. Best would be if in some manner this area is conical, since than the view would not need to be exactly on the object. Another, in our opinion better, approach is to create an area around the selectable object which can intersect with *v*. With a cone there would be a maximum distance before the cone would cover everything in the field of view. In this manner this is not a problem.

An alternative would be using a cursor, however this would mean additional cognitive load on the user as he would have to both control a cursor and the viewpoint. Since opinion was split on this issue we opted for the conceptually easier 'viewpoint selection technique' which is also often used in VR applications, especially immersive applications.

5.3.4.2 Menu Selection

Now we can start defining the menu structure that is designed to support menu selection. Another constraint we have put on this menu structure is that it can somehow show relations between items. Menus usually do this already in one dimension, but we would like it to be possible for at least another dimension, making it possible to browse through two dimensional data structures and visualizing them. Better even would be to make an n-dimensional selection tool but this is a complicated manner and it is very questionable if the average user would conceptually grasp this.

Yet another constraint is usability. The option that is selected and the associated 'activation' of that option should be easy to grasp, hence the results from the initial user study have been used to create a mechanism that affords reaching out to the selected option to 'activate' it. An immediate result of this is that the selected option is closer to the user than the others, one could say 'within grasping distance'. We will now describe the specifics of the selection system we designed: The Selection Wheel.

As visible in figure 20 the user is presented with a wheel that can pop up around an object (though the object is not shown in the rendering) as soon as the relevant object is selected. The option closest to the user is the currently selected option, which should be made obvious by visual feedback like making it slightly larger or highlighting it. By executing the 'select action' associated with the interface device one then selects this option. Moving to another option is done by doing the associated action browsing actions for each device. For gestures this would mean that if one were to wave up, the wheel would rotate clockwise, and counter-clockwise for a downward movement. This would happen in a discrete manner, the menu would rotate to the next sphere up- or downwards. This might later be extended to allow faster browsing (for example, skipping a couple of spheres using 'double-clicking' or more powerful gestures). However it seems important the mechanism remains discrete as the system will not interpret user's select actions wrongly and the status of the menu is also always communicated clearly to the user (imagine a pull-down menu where you could be between choices).

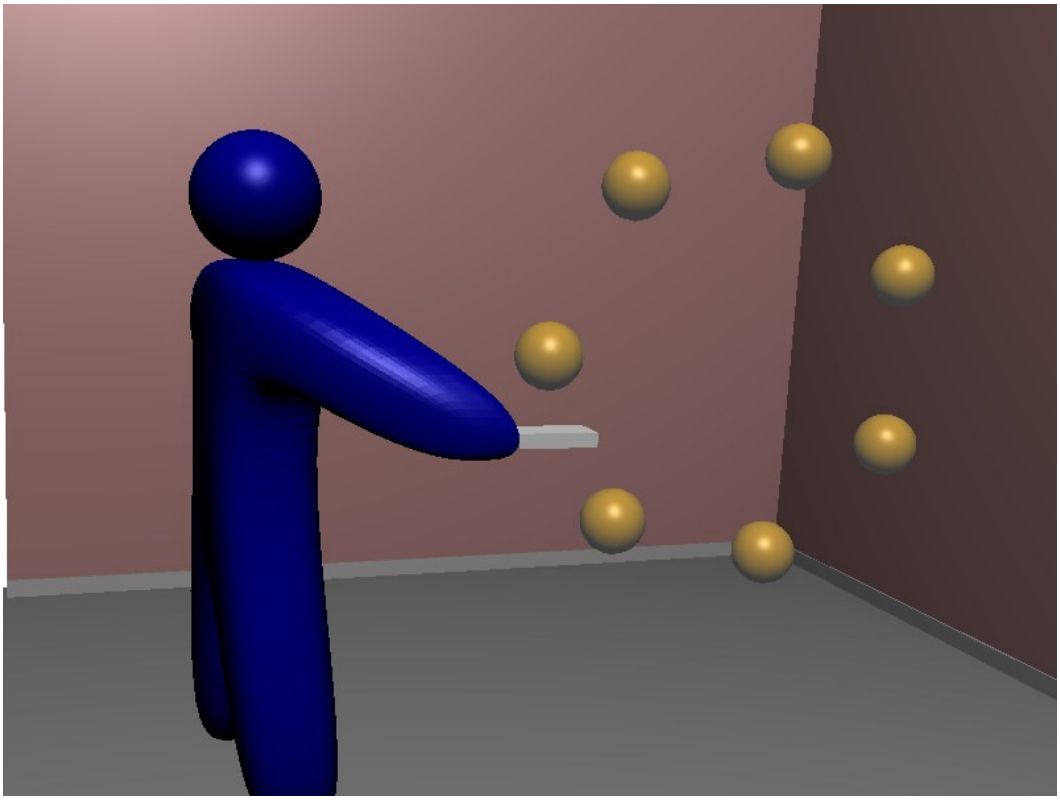

FIGURE 20 - THE SELECTION WHEEL: THE ORANGE GLOBES ARE PROJECTED IN THE VIRTUAL SPACE

By putting the selected 'node' closest to the user it is hoped to make use of the third dimension available in virtual reality in an obvious way; by making it a wheel it is hoped the user will intuitively understand that this wheel can be rotated somehow to arrive at the next or previous option and thus relations are visualized. Furthermore we have designed a number of animations as the wheel is opened. The nodes in the wheel will animate outwards from the center to reinforce this 3-dimensional sense of the wheel as well as a mechanism that shows when a new wheel is opened. As can be observed the wheel in figure 20 is put in a vertical plane, of course a horizontal plane could also be chosen, depending on screen dimensions and the likes. In this case it would rotate much like a tire rotates.

We will now show how a wheel in the horizontal plane is constructed, using the 'navigation plane' as described earlier. Because of this we have an easy time calculating the *x, y* and *z* coordinates of the globes (since they are in a plane and we are showing calculations for a horizontal wheel the z-coordinate will actually be a constant).

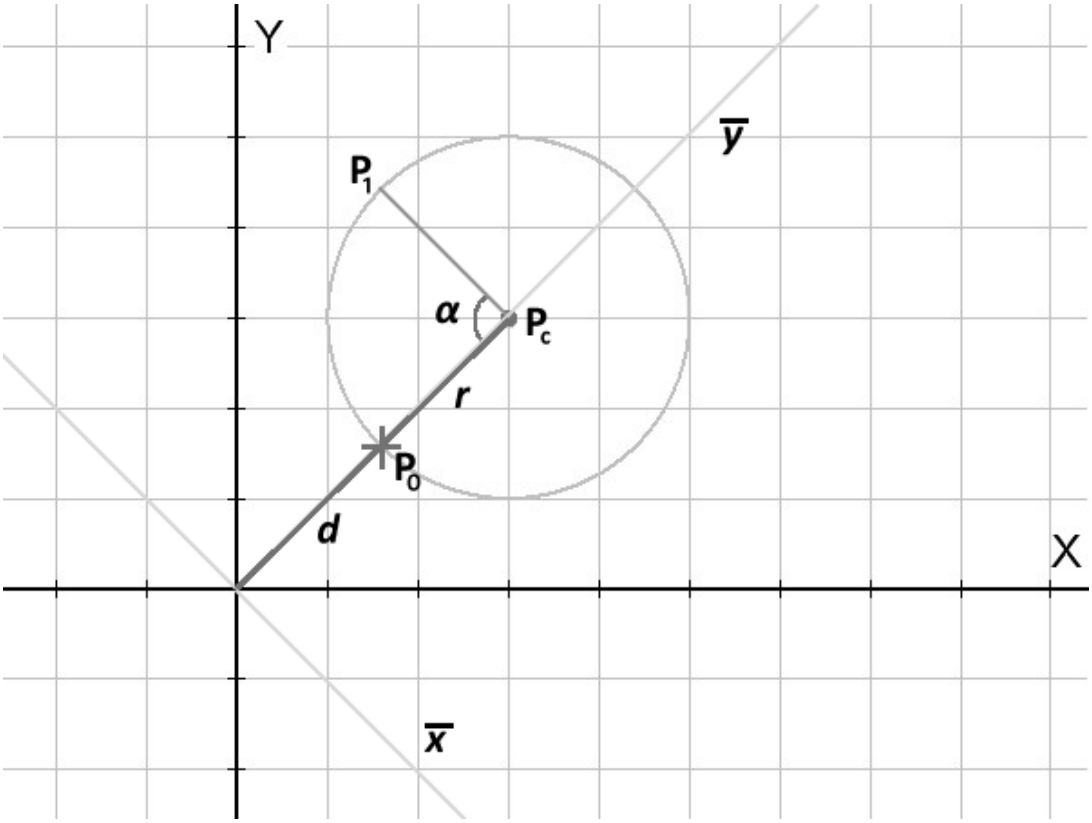

FIGURE 21: AN EXAMPLE OF THE CONSTRUCTED POINTS P_0 AND P_1 FOR CASE N=4

Now as you can see in the picture above there is an absolute coordinate system with an X and a Y-axis. Since we have the relative coordinate system we can define two new axes which we shall call the $\bar{x}$ and the $\bar{y}$ axis. Now if we want to put a wheel at a distance of *d* units from the user, and have a radius *r* we can define the points P_c , the center of the circle, and the point P_0 , the position of the currently selected sphere, in terms of $(\bar{x}, \bar{y})$ coordinates as follows:

$$P_c = (0, d + r)$$

$$P_0 = (0, d)$$

Now if the selection wheel has *N* objects in the circle we can define the position P_n , being the position of the n-th sphere in the circle, as follows (once again using $(\bar{x}, \bar{y})$ coordinates and assuming the y-axis is the axis pointing forward)

$$P_n = (\sin\left(n \cdot \frac{360}{N}\right) \cdot r, d + r - \cos\left(n \cdot \frac{360}{N}\right) \cdot r)$$

This is basic goniometric math using the unity circle, but is included for completeness sake. Notice how this also works for P_0. The reason we are able to use this basic math and do not have more complicated calculations to do is because we are using the framework laid down for navigation, and placing it in the relative coordinate system,

instead of calculating the 'real world position' of these items. Depending on the implementation a translation between these relative coordinates and the world coordinates might still have to be made, but that is an implementation issue.

Now that we have the position of an object in the wheel we can place these objects at the correct place. If we now wish to rotate this wheel we will have to do that over an angle of $\frac{360}{N}$ degrees (in the example above in figure 21 it is $\frac{360}{4} = 90°$) to either the right or the left. Certain implementations might have rotate actions for objects included; if this is not the case we would suggest using a standard rotation matrix.

The use of a wheel also makes possible showing two dimensional related data. One could combine both a horizontal and a vertical way to browse through it. Data related in more dimensions could even be shown by adding more wheels (diagonally oriented for example) however this might swamp the virtual space with (in this case) spheres and it remains to be seen if that is good. Two dimensions however are also quite well supported and a user should be able to see how something is structured. Since this wheel now supports both menu operations and multidimensional navigation our requirements on the selection mechanism are fulfilled.

In immersive environments such as stereoscopic displays or HMD's it is expected that this design will work quite well, since one can 'feel' the wheel in space and get a spatial sense of the wheel without even moving it. A gesture interface might also prove very effective here, users did indicate they would feel like reaching out to close objects that look menu related, and in this subconscious gesture the key to make the interface very natural might be found.

5.3.5 Mapping the input device

Since we have now defined what is possible in the application we can start mapping the input device in an understandable way to the tasks which are possible. We will shortly list these tasks first and then explain how and why the WiiMote and the SpaceBall are mapped to these tasks.

5.3.5.1 Basic Interaction

In our application we can define 7 basic interactions in the interface (more complex tasks can be done using these interactions e.g. 'entering the historical simulation' could be done by selecting a certain object, selecting it in a menu, etc., hence these interactions are defines as the *basic* interactions)

1. Moving the viewpoint
2. Orientating the viewpoint
3. Selecting
4. Unselecting
5. Opening a menu/selection wheel
6. Closing a menu/selection wheel
7. Navigating the menu/selection wheel

Some of these interactions could be divided into even smaller parts, for example moving the viewpoint could also be 'moving forward', 'moving backward' and 'moving sideways'. We choose not to do this since on an interface device these tasks should be controlled with one mechanic. Often even moving and orienting is done with one mechanic. Since these interactions are defined with the purpose of mapping devices on them we deemed it unnecessary and will use these as our 'basic interactions'.

5.3.5.2 WiiMote

We have used the WiiMote together with the nunchuk expansion. The reason for this is the joystick offers a good control for 2DOF control over the movement due to its isometric nature (Bowman, Kruijff, LaViola, & Poupyrev, 2004) and is seen as a natural interaction device, operable even by people with a brain injury (Wallergård, 2007) (Mine, 1995). Two handed interaction has also been shown to be very natural to people, assuming off-hand actions are kept simple. What is usual for persons is to use one hand in support of the other. This means that if a user is trying to accomplish a goal, the offhand should support this task, without requiring intricate collaboration with the main hand. In our terms this would mean not mapping navigation and selection to one hand, or to 'both hands' but rather one task on the offhand and another task on the main hand. (Verheijen, 2004)

In the case of the WiiMote and nunchuk together this would mean using the isometric controls of the nunchuk for navigation and the isomorphic controls (gestures) of the WiiMote for selection. Hence we have used the joystick on the nunchuk for basic 2DOF movement which already has been shown to be natural to people. Using a 'modifier-button' on the nunchuk we have defined vertical and sideways movement for persons. This means that with the nunchuk the user is able to move along all 3 axes and rotate over the z axis (yaw). This should be enough for basic 'moving' of the viewpoint. Orienting the viewpoint is done using the gesture sensitive WiiMote. The idea behind this is that one can move the viewpoint with the offhand, while doing a more complicated task (orienting the viewpoint) with the main hand. A user will look around with the WiiMote as if the virtual head is mapped to the movement of the WiiMote. Hence if you point the WiiMote up at a certain angle, the user's viewpoint will be looking up at that angle as well. To make sure a user does not inadvertently orientate the viewpoint if he is gesturing with the WiiMote a button-combination will have to be pressed to activate the orientation function.

Walking itself is chosen to be at one speed. This is chosen since people often also walk at one speed. There is no conscious thought of 'I will walk at 80% of my speed over there', only a simple 'walk or not walk' hence we have used that mode of interaction. The walk speed in the application is mapped such that it is a 'normal' walking speed (between 1 and 2 m/s).

In bowman's taxonomy this would be defined as follows: We are using gaze-directed steering with a constant velocity, operated by constant input.

As for selection we played around with the idea of gestures, however making a 'grab' gesture with a WiiMote does not seem like the best idea (it affords pointing, but there is really no way to differentiate between normal pointing to an onlooker and a selection point), and since users seemed very comfortable using buttons we have used a button press to select. The button we used is the A-Button which is also the button used in every WiiMote application to select. Conversely, using the same logic, the B-Button is used to unselect.

Operating the selection wheel is done using gestures because of the amount of fun and interactiveness with the world. The wheel shape is used as a basis for the gesture. The gesture used is as if the person is giving the wheel a spin (comparable to 'wheel of fortune'), hence with a swing to the left (in the case of a wheel in the horizontal plane) the wheel will turn clockwise (if observed from above with the user standing south of the wheel). To open and close the selection wheel we once again use the A and the B button for the sake of consistency. This does mean the software will have to keep track if a wheel is open, since it will have to respond correctly to button presses.

5.3.5.3 SPACEBALL

For the SpaceBall we have decided to also use gaze directed steering with a constant input and constant velocity. The SpaceBall does not have the gesture capability of the WiiMote so some things had to be done differently though.

We will start with the description of moving and orienting the viewpoint. This is done by mapping movements made with the SpaceBall directly to movements made by the viewpoint. If one rolls the SpaceBall left, the viewpoint will roll left, if one moves the SpaceBall forward, the viewpoint will move forward, etc. We believe this will provide a user with 6DOF in an incredibly intuitive way. The rolling movement was included here even though it was not in the original design of the navigation, since it felt odd to leave it out. Rolling using the SpaceBall is no more difficult than moving forward.

To select and unselect we use buttons on the SpaceBall similar to a mouse. The 'left button' is used as a select button while the right button is used as an unselect button. We use these similar to those on the WiiMote so they also serve as an open menu/close menu button. By doing this mapping of buttons we have essentially created an easy way to browse through multiple levels of menus. By selecting one always moves a level up, by unselecting one moves down a level, until one unselects the object. This argument also holds for the WiiMote.

To rotate the wheel one simply rotates the SpaceBall like one would want to rotate the wheel. This means multi-dimensional wheels are also possible up to a certain extent. Since this operation is also used to move around it is implicitly not possible to move around while the wheel is open with the SpaceBall. This has not been included in the design but if one would use the Wheel in another application a solution will be necessary for this.

5.3.5.4 ON SCREEN HINTS

As we already mentioned no matter how you map the devices, it remains a very subjective experience which often yields an unclear ranking between devices. Since we are trying to test the best method of interface another thing we will be testing if performance (and thus the overall experience) in this Virtual Reality museum can be improved using other methods than just the interface device. One such method is the use of onscreen hints.

By providing users with on screen hints on how to operate the devices, the inherently unnatural element of said devices (a joystick or a SpaceBall are after all interface devices and not natural objects) is overcome. Even if an implementation or mapping of a certain device is not obvious to the user, the on-screen hints will explain it. Since our navigation and selection is quite basic in nature, this might let the user learn how to operate the device more smoothly without reading heavy manuals.

A drawback could be that it destroys the sense of immersion which is one of the 'fun' things of the VR museum experience since these hints will remind the user of the fact that he is using a computer interface. This was also researched during the user study.

During the design of the hints we have tried to focus on simplicity and on keeping them small. This means that we have tried to put most of the interaction possibilities on one picture. In the case of the WiiMote we opted for two pictures since you also use two hands and in a way, two devices. Furthermore in our design only relevant on-screen hints are shown. As mentioned you cannot move if the selection wheel is open, hence on screen hints about moving your viewpoint are not shown when the selection wheel is open and conversely, if there is no selection wheel open on screen hints about operating the selection wheel are not to be shown. We will now present the collection of on screen hints designed to be used in the application.

The SpaceBall:

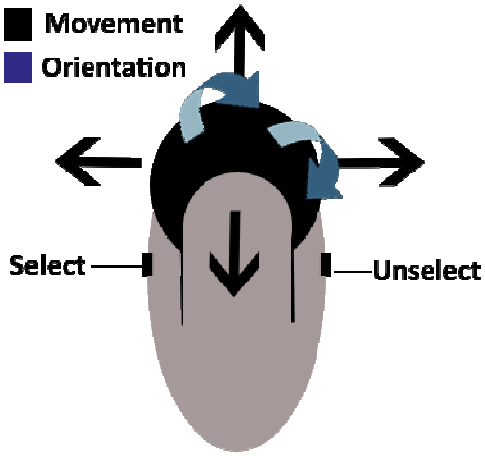

FIGURE 22: MANIPULATION OF THE VIEWPOINT

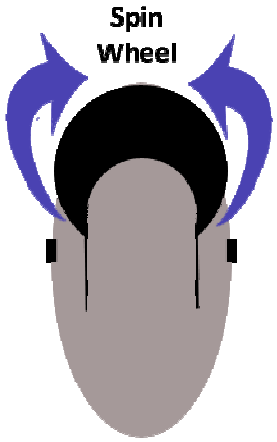

FIGURE 23: OPERATING THE SELECTION WHEEL

The WiiMote:

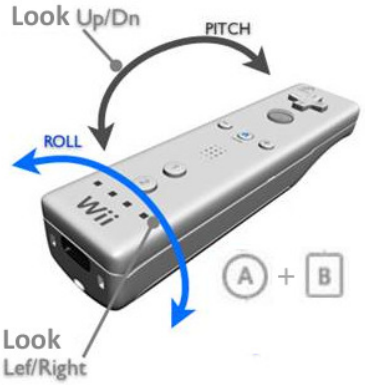

FIGURE 24: ORIENTING THE VIEWPOINT

FIGURE 25: MOVING THE VIEWPOINT

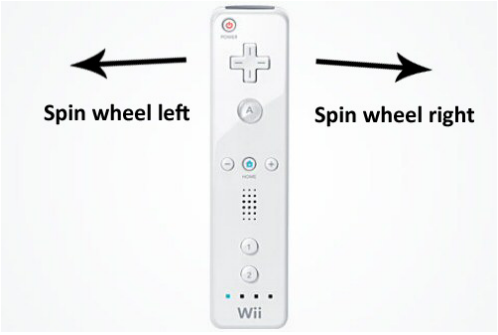

FIGURE 26: OPERATING THE SELECTION WHEEL

Since the WiiMote has very recognizable buttons we have select and unselect hints that are only activated when one is able to select an object or unselect an object:

FIGURE 27: SELECT AND UNSELECT HINTS

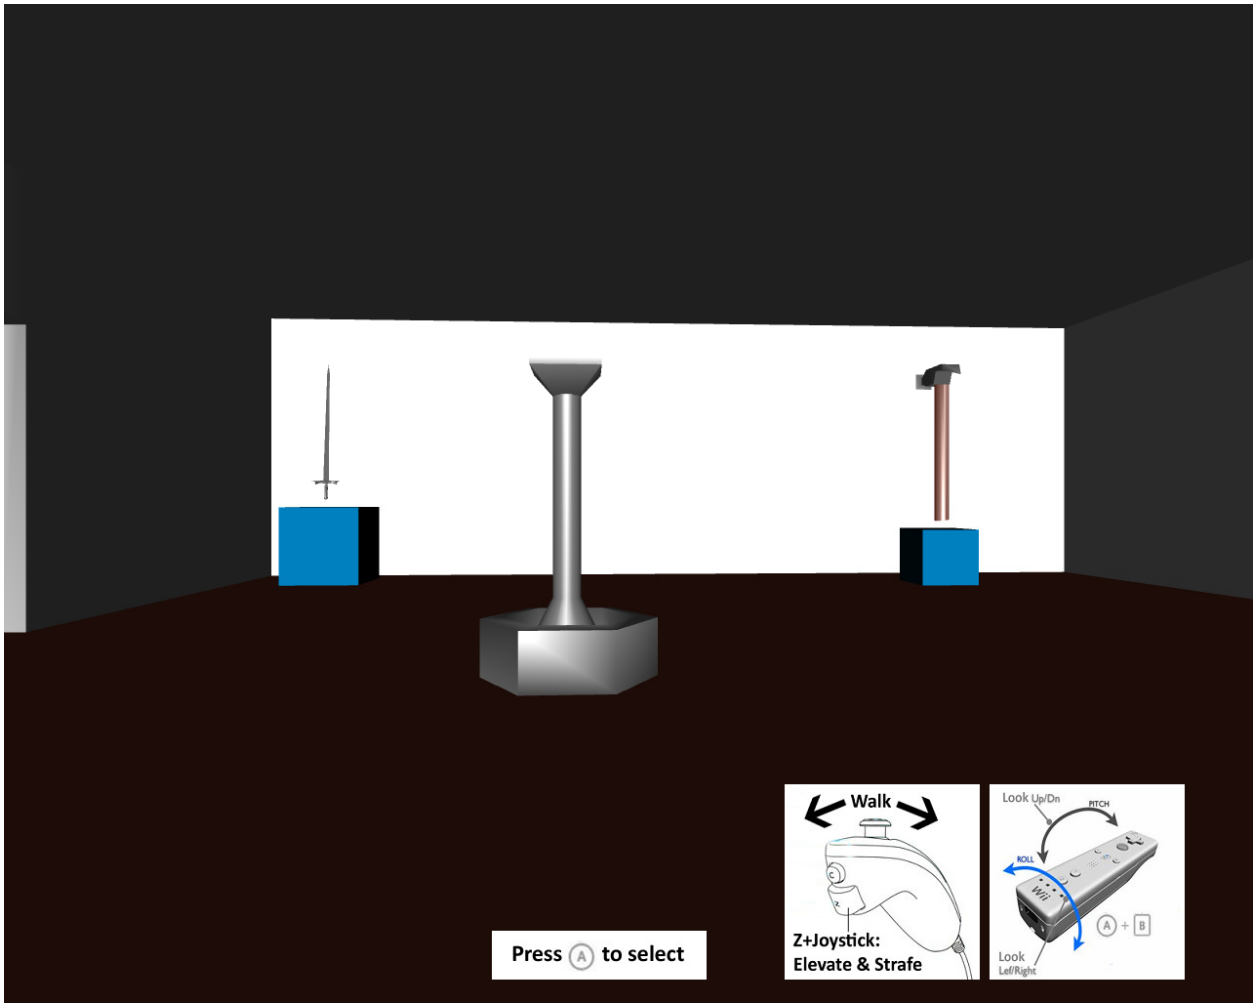

FIGURE 28 - ON SCREEN HINTS FOR THE WIIMOTE USING THE APPLICATION

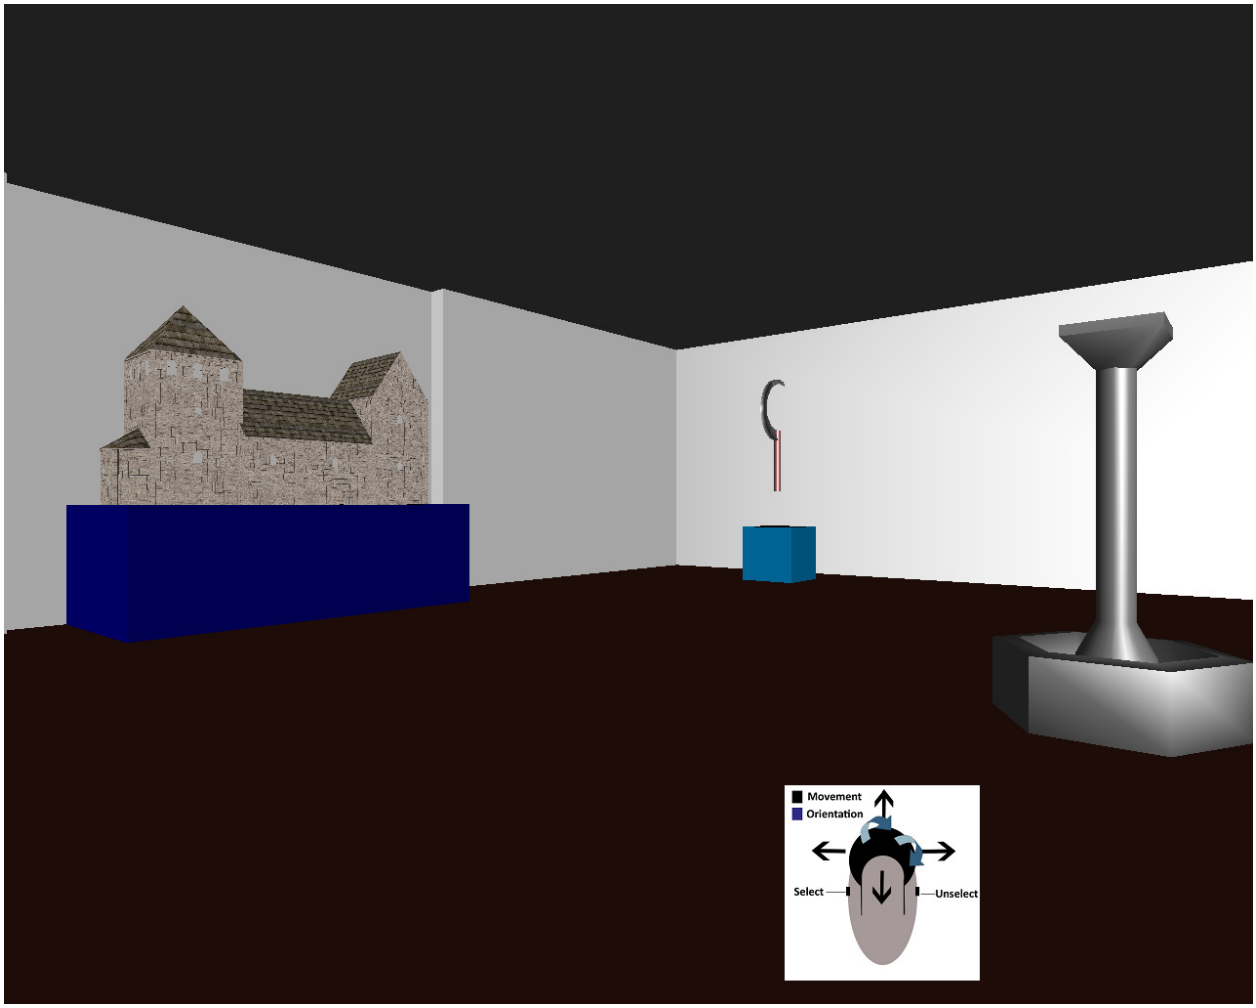

FIGURE 29 - ON SCREEN HINTS FOR THE SPACEBALL

As can be observed we used figures looking like the real objects. Furthermore the hints for the WiiMote are based on those as seen in many games as is the wording (elevate and strafe as opposed to move up or sidestep). For the SpaceBall we chose a 2-dimensional representation and trust that this will be enough to make sure the user also understands a 3rd dimension can be used. It will be apparent quickly enough once the user since the SpaceBall does afford using it in every direction.

5.4 Overview

To create the design we have started out with an initial user test. Using these results we have defined the essential task scenarios that can provide a good experience as users expect it from a VR museum. Navigation through viewpoint manipulation and selection of information on exhibits turned out to be the most important tasks. A much appreciated feature was the inclusion of a historical simulation.

Using this data we have created a design where a user will be able to navigate through an environment based upon the Kulturen museum in Lund, where exhibits are shown as in a regular museum and the user is able to navigate using multiple degrees of freedom. The user is also able to access a historical simulation from this level. To navigate through information we have designed a menu system that can be extended to allow multiple dimensions of data to be browsed; this mechanism is called the selection wheel.

We have then taken the atomic tasks of this design and mapped functionality of the WiiMote and the SpaceBall onto them in what we thought was a natural way. To accommodate novel users we have included on screen hints to see if they can improve user performance, and thus the user's experience.

Now that we have a design we will tell some more about how we implemented this design in the next chapter, including some screenshots of the actual application in action.

6 THE APPLICATION

In this chapter we will give a technical description of the application/framework based on the design in the previous chapter. This framework was developed and used for the user study described in chapter eight, nine and ten. We will show how we implemented certain design decisions, which software was used, how the software is constructed and some hands-on documentation which might prove useful if one would like to use the software we created. If you are interested in the research process this chapter is not the most interesting, however as for practical tips on the execution of certain paradigms, or if you are interested in an example of VR development, this is probably the place to be.

6.1 SOFTWARE USED

One of the first choices we had to make was the choice of software. There are many visualization packages and libraries available for as many languages. Our project had a few specific requirements to which we tailored the choice of software:

- Generality: The hardware platform on which this application would be used was not known beforehand, hence it had to be easily adaptable to CUBE systems, projectors or desktop computing.
- Interaction: The focus of our project is on interaction, not quality of graphics or speed of the engine. For this same reason libraries focusing on physics or manipulation of 3D objects were also rejected, first and foremost was basic interaction.
- Simplicity: The relative inexperience of the student performing in this project in the field of 3D graphics and the focus on other steps in the development process, specifically both user tests, meant the implementation had to be kept relatively simple hence a package or library with a very steep learning curve would be rejected, this often concerned graphics and adapting them to many displays.

Based upon these arguments the software used for most of the implementation and programming was EON Studio[12]. This software had its own graphics engine and the ability to import the most common 3D model formats. Furthermore it had a large focus on interaction through a graphical authoring tool using interaction 'nodes'. A lot of pre-programmed content was available while one could extend this content with own scripted events that resemble objects in common OO languages (though they cannot be instantiated automatically). Using these features together provided us with easy-to-access interaction events that could be related to any 3D model, on which various exterior modifiers could be applied as well. Another added factor of using a more simple approach such as this is that it might be easier for people without programming experience to work with the provided framework to design a better museum, and immediately integrate this work into the framework.

Another advantage of EON was its specialization for VR displays. One can alter the render-engine in such a way that it uses either a stereo display, a head mounted display, a desktop display, I-Glasses or many more options. This way no extra work was needed to adapt our implementation to different platforms and thus time could be saved on this.

Alternatives considered were XNA by Microsoft[13], OpenSceneGraph[14] and JMonkey[15]. However all of these libraries, while having some focus on interaction, provided very little in the way of rendering possibilities for exotic

[12] http://www.eonreality.com/

[13] http://www.xna.com/

displays, or were lacking in focus on interaction. The focus here was mostly on providing fast, high-fidelity graphics and including a physics engine for realistic animations.

Software was modeled in Blender[16]. This was exclusively chosen for its open source format, hence it is free 3D editing software which is also widely supported and has an active userbase. It can export to every popular 3D format. The choice here was less important though since many good 3D modeling tools exist nowadays, with similar interfaces.

6.2 3D Models

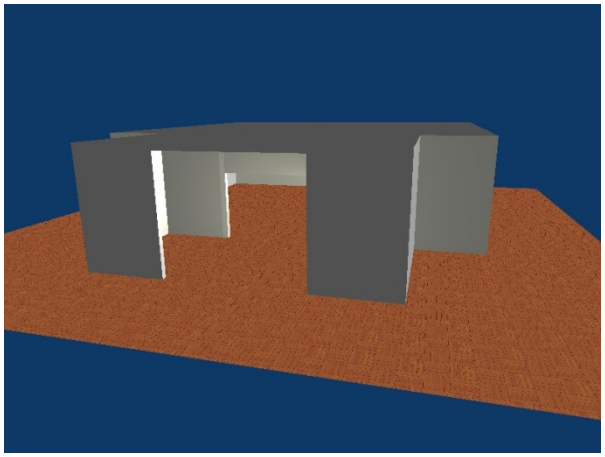

FIGURE 30 - THE VIRTUAL MUSEUM ROOM

We created a few 3D models based upon objects found in the Kulturen museum to provide the user with an authentic museum-feeling during the user study. The objects were photographed and simple low-polygon 3D models based upon these photographs were made. We settled on using six models in our simulation room, which should provide the user with enough choice, and enabling us to include a search task in the user study.

The actual room in which it was placed was modeled to resemble a normal museum room without too much clutter. It had two sections so search tasks could be made more challenging by placing objects around the corner. The space used was considered to be quite normal for a museum room, based upon personal experiences.

The 3D models of the exhibits were not rich in detail, as we focused on the interaction in this project. The user study did take this into account and feedback was asked about the required level of detail for an application like this. The objects modeled in the application were: A medieval longsword, a hammer used to create wooden shoes, a cooking pot on a cooking ring, a sickle, a herb garden fountain and the Drotten Church. To indicate the level of detail used we will show the fountain and the cooking pot as rendered in EON and the photos they were based on.

[14] http://www.openscenegraph.org/projects/osg

[15] http://www.jmonkeyengine.com/

[16] http://www.blender.org/

FIGURE 31 – MODELS AND THE PHOTOS THEY ARE BASED ON

As can be noticed there are no textures or bump mapping or any other sort of mapping. Colors are used though in the case of objects that had a wooden handle. This was deemed enough for this framework as objects are quite recognizable. It will probably be the case that in a finished application, more detail is called for.

6.3 A Technical Model of the Application

The first model is a state chart that shows how the program evolves through different states and which actions are available in each of these states. There are three basic states that can be identified:

I. Navigation State
In this state the user will be able to navigate through the environment. He will be able to orientate himself and move, but this will not change the state. If he selects an object however the user will enter a Selected State

II. Selected State
In this state the user will have an object selected. This means that the selection wheel is visible and in this way the user can interact with the object. The user should not be able to orientate and move himself. The user can unselect however to go back to the navigation state. By pressing select the user will run a command that is associated with the currently selected wheel option. By rotating the wheel the associated Action Command will change

III. Action Command
In this state the user will run an action that is associated with the action command that was just selected. This can be anything such as opening a new wheel, opening a video or entering another simulation. It is also possible to stack multiple instances of this. For example one can open a new wheel, from where one can run other actioncommands, thus entering the state associated with those actioncommands. The idea is that the structure is always the same as shown in figure 32, which guarantees that the select and unselect button will always move between the current 'actioncommand' state and the state where a user came from thus providing consistent menu navigation.

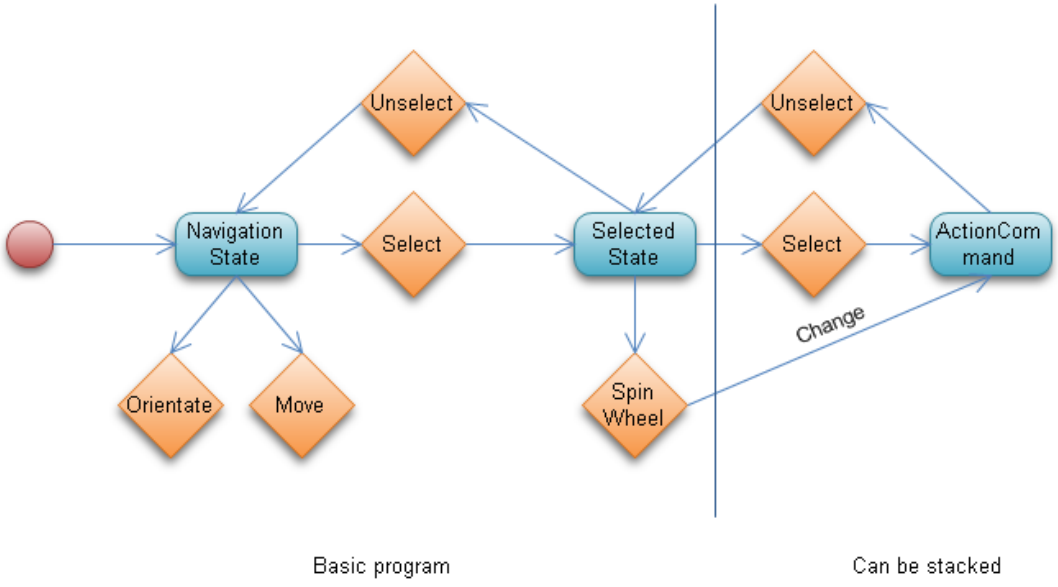

FIGURE 32 - POSSIBLE ACTIONS IN THE PROGRAM AND THE WAY THEY CAN BE STACKED

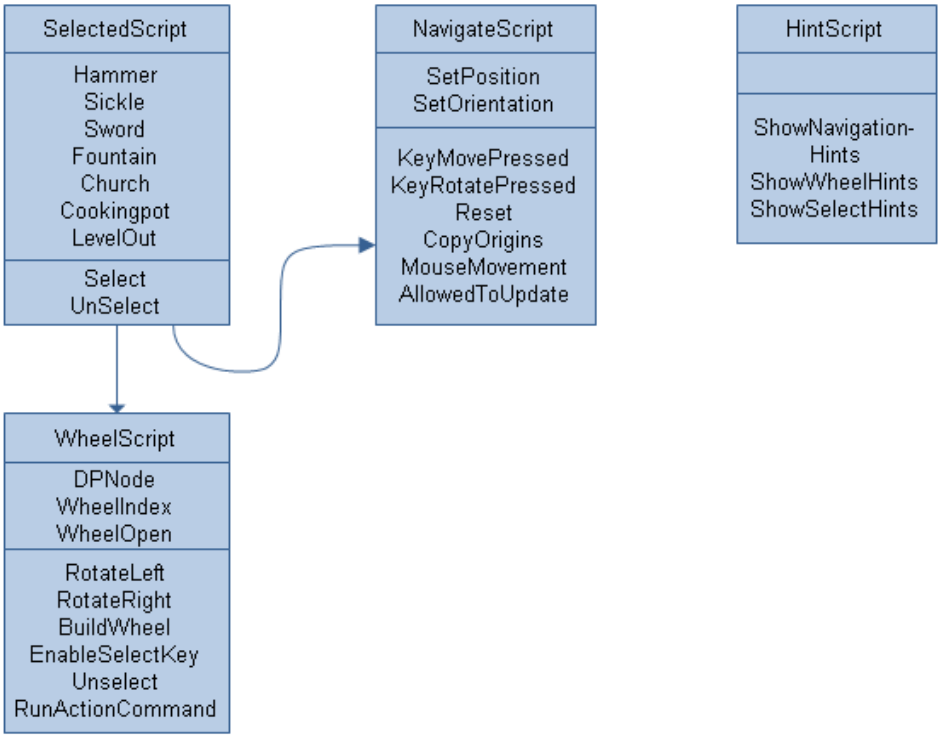

FIGURE 33 - A SIMPLE UML DIAGRAM OF THE SOFTWARE AS IMPLEMENTED IN EON STUDIO

The figure above shows the setup of the scripts in EON. It is modeled as a UML diagram but please keep in mind that EON uses JavaScript which is not an object-oriented language. This is also the reason why so many variables (such as the exhibit-specific fields for selected objects) are 'hard-coded' instead of instantiated objects.

SelectedScript takes care of selection and unselection and contains a Boolean field for every exhibit indicating if it is currently selected or not, so this information is available for any other script. WheelScript takes care of the generation of the wheel, keeping track of the currently selected option in the wheel, rotating it and starting whatever the user wants to start when he selects an option in the selection wheel. NavigateScript takes care of the synchronization used for the relative system of coordinates when input buttons are pressed. HintScript is only used in the version with on screen hints. It has internal code for showing hints specific to a device, however it is advised to use these general functions, to make it easier to switch between devices.

The arrow coming from the SelectedScript towards the NavigateScript indicate how when selection is done and thus the user is animated to the selected object (the zoomback technique) the SelectedScript notifies the NavigateScript about the position changes to make sure the relative system of coordinates and its synchronization with the camera remains intact. The other arrow indicates the forwarding of select and unselect presses (which are handled by the SelectedScript) towards the WheelScript in the case a selection wheel is open.

To try and make all of the above somewhat more visible for the user we will show a quick glance at the node screen of EON below. Please note how you can see SelectedScript, RotateScript and DOF_PositionScript outlined in orange (which correspond to SelectedScript, WheelScript and NavigateScript in the UML diagram).

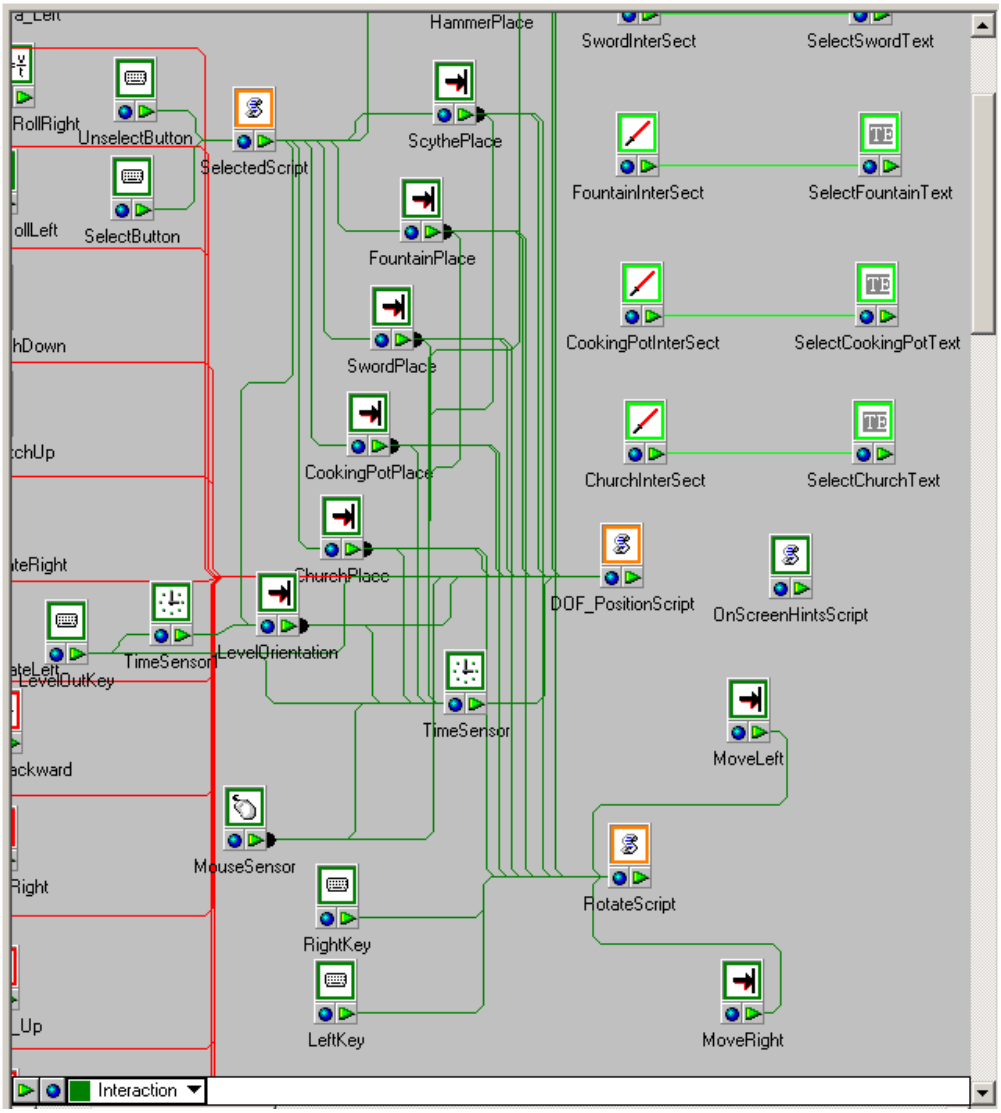

FIGURE 34 - THE NODEVIEW IN EON

6.4 From Design to Application

We will try and give a short overview of the most important matters in implementing our navigation system as well as the selection wheel. This is not however a detailed explanation on how EON works or how it is programmed line by line.

The navigation system as described in chapter five was comprised of a relative system of coordinates in which the user moved the camera. By keeping the origin of this relative system synchronized with the position of the camera (and thus making sure the camera is at position (0,0,0) in this relative system) we had the advantage of easily controlling the amounts of DOF controlled as well as much control over how we mapped the interaction. In EON this translated to using a 'DOF' node in the simulation tree under which the Camera node was placed at position (0,0,0). By making sure that every time the camera was moved or reoriented the position and the orientation of the DOF node was also updated, and resetting the position and orientation of the camera to zero, we had in effect implemented this method.

To implement the selection wheel we used a technique that is quite new in EON which is editing the simulation tree run-time. When a wheel is created (done by the user selecting an object) the application would find an associated text file and use this to generate a wheel. In this way it will hopefully be easy to extend the application in the future. Using this text file the number of spheres in the wheel is determined after which all the spheres are put at their right spots within the DOF node through an animation (to give the user an extra depth cue), thus making sure we don't need to do any more translations or rotations on them to align them with the users viewpoint. This means that while the simulation is running, spheres are being copied under the DOF node and removed again when one closes the wheel. For more information on this (still) exotic procedure we would refer the reader to the Scripting reference of the EON User Manual included with the software.

6.5 Using the Interface Devices

To make sure our application mapped the basic tasks as defined in 5.3.5 to the relevant interface devices we chose an approach that allowed us to accomplish a lot without extensive programming. By using an application called 'GlovePIE'[17] we were able to map interactions with exotic interaction devices (such as the WiiMote and the SpaceBall, the GlovePIE software allows for many more) to keyboard and mouse action. It is relatively simple to read mouse and keyboard actions in EON itself and use them to trigger certain events. Now all we had to do in EON is make sure the program was able to do the basic tasks in a correct way using the mouse and the keyboard. Then, using the scripting capability that is provided by GlovePIE we could map these actions to keyboard and mouse movements.

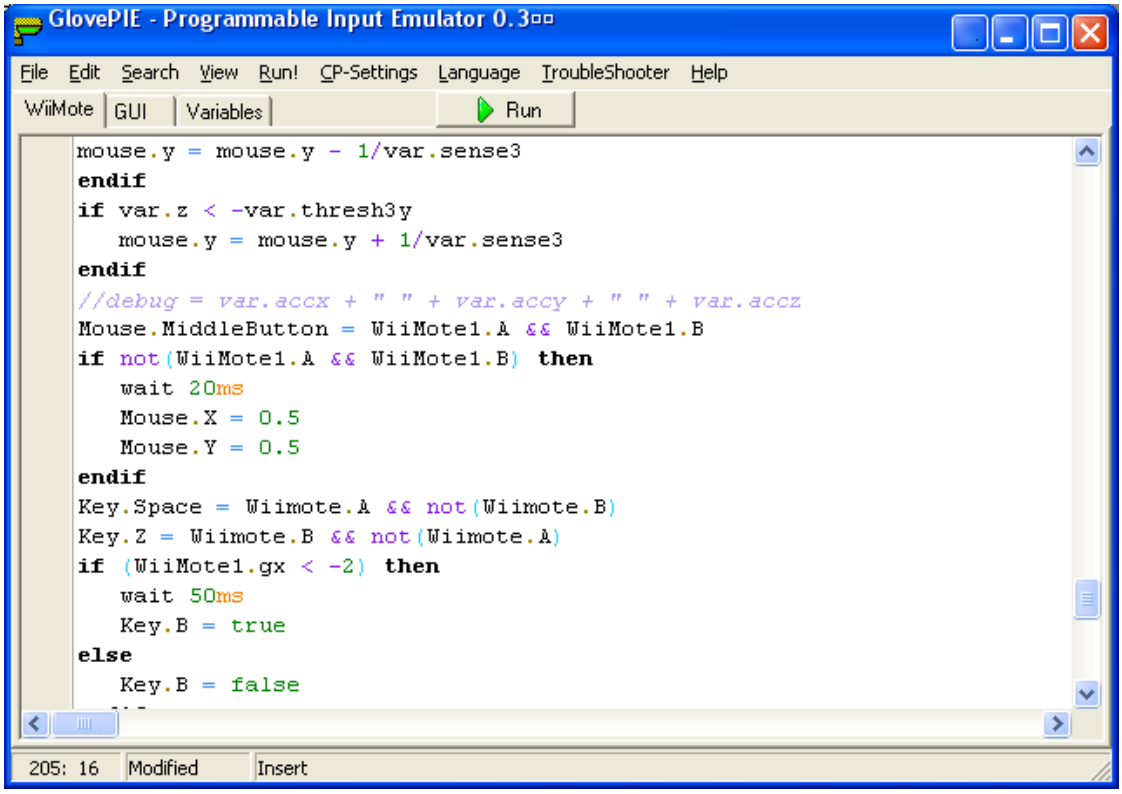

FIGURE 35 - EXAMPLE OF THE GLOVEPIE SCRIPT USED FOR MAPPING THE WIIMOTE

[17] http://carl.kenner.googlepages.com/glovepie as visited during the entire project

This approach has several advantages. First of all, the inherent scripting capability in GlovePIE is quite extensive and designed to address real-time constraints as well (though not on a precision level required by normal control systems, since it is still a script running in the Windows Operating system) through the support of semaphores and wait operations. It remains easy to understand though and as such the interaction with the application and the manner in which the device operates can be changed or edited without knowing anything about how EON works, as long as one knows to which keyboard presses and mouse movements one has to map. This will make configuring and fine-tuning for specific devices a lot easier.

Then there is the fact that GlovePIE is OpenSource and thus one can edit the interaction part of the application without having to buy the expensive EON editing software, instead just using the EON viewer. This might make it more manageable for a real museum to use.

Furthermore, using this setup it is easy to add new interaction devices to those already supported. Since the implementation made in EON is (supposedly) already finished and well-designed, one only has to make a script that maps the input of any interaction device to that of EON. Nothing more than GlovePIE is necessary. This layered setup is shown in the figure below.

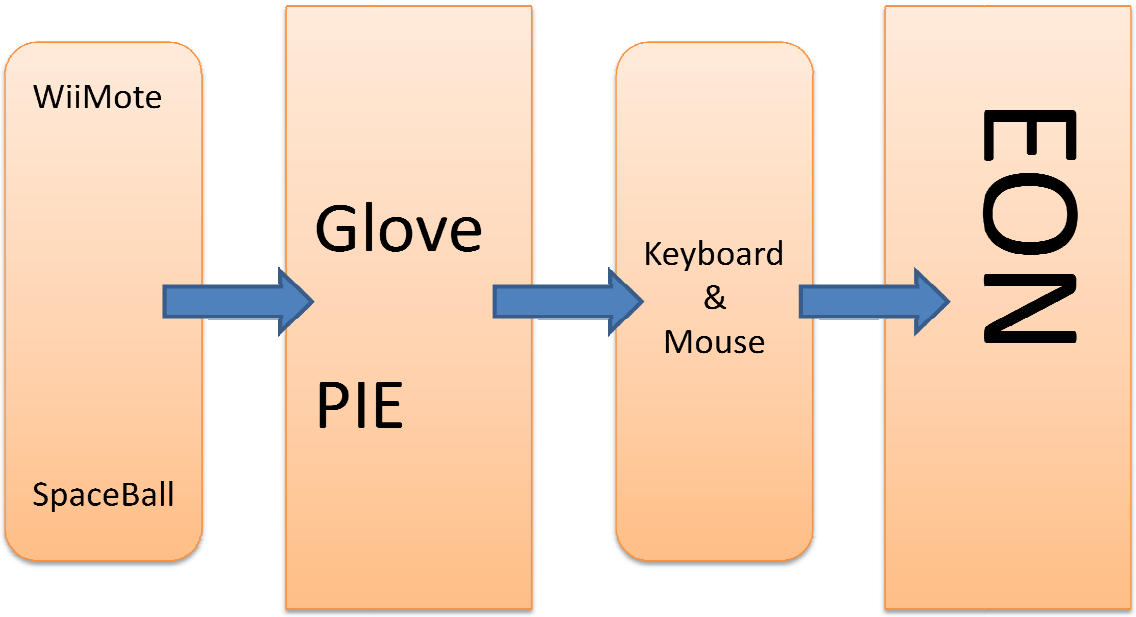

FIGURE 36 - RELATIONSHIP BETWEEN GLOVEPIE AND EON

As can be observed, by simply adding devices to the 'device space' before GlovePIE and then creating a script in GlovePIE, the right part of the diagram can remain unchanged, thus allowing for any device to be added later in a relatively simple way.

There is one drawback, and that is that the EON system cannot interact with the interaction device. While this is not important for interaction devices that have no feedback mechanisms, this is a drawback for something like the WiiMote which can give audio, haptic and visual feedback. There is a workaround, where EON can simulate mouse presses in certain patterns, which in turn can be read by GlovePIE and then translated to feedback signals. This is however a bit contrived, and if feedback through the device is an important part of the design of your application we would recommend using the EON SDK to write an EON node for the device; keep in mind that this is not a trivial task and may cost quite some time.

6.6 ADAPTING TO DIFFERENT DISPLAYS

Adapting to different displays is quite easily done in EON, however there are a few things you should take care of. For example in our case, we mapped the application to the 'podium' which is a setup that contains three screens that are positioned at a certain angle of each other. In EON this means that the used camera node has to contain three other camera nodes at these same angles as those used in the podium with a correct field of view. Each of these cameras then map to the three available viewports which are then aligned. In effect this means if one moves the 'master' camera, all three viewports are updated. While this is not much work one should still take care that this is correctly done for every display you use, and this example is included so that users may consider it when expanding this application. Remember to make it scalable to multiple viewports.

7 USER STUDY

The framework described in the previous chapter has been used to answer certain questions we have about virtual museums. We would like to make sure the user has a good experience there, that he would like to repeat, and we would like the interaction to be as easy as possible in this context. If this is all successful we could use the results to bring the past to life using these modern techniques and use them to educate people about history, which in itself is important enough.

However to determine if this framework is successful in the described goals we have tested the user experience and to that end there we planned a user study after the implementation was done. Using this study we gathered data that can confirm certain hypotheses. Furthermore the results of the user study can be used for further development of this application and might perhaps give rise to replicating certain approaches taken in this project to other projects.

7.1 OBJECTIVE

During the test the objective is to gather both qualitative as quantitative data. It is believed that getting only one type of data will either lead to vague conclusions, in the case of qualitative data, or not be complete enough in the case of only quantitative data because the user's experience is very important and this is not something which is easily quantifiable. Since we are doing an Application Specific test that requires users there are specific tasks that the user needs to do that need to be formatively evaluated. It is researched how well the interface supports these tasks and if the experience is what is expected.

To this end the test contained three parts. First there is a small interview, where the user is asked about his or her experience with 3D navigation and the WiiMote. The subjects were then asked what they expected from a virtual museum regarding the level of interaction, the amount of fun and the educational value. Secondly the users had to perform a set of tasks that could be typical task in this design of a virtual museum. These tasks are based upon the user task scenarios described in chapter five. During these tasks a number of metrics were measured to see how well the task is performed. The users were asked to 'think aloud' while they are doing this to document design flaws and to try and get an unbiased view of why the user is taking certain actions. After the test the users had to fill out a small questionnaire to get quantitative data about the experience. Thirdly there was a post-test discussion where the users were asked open questions about the experience, and if it could add something to a real museum visit.

There are two questions that we hoped to answer using the data collected in these tests:

1. *Do on-screen cues significantly improve user performance in a virtual museum, regardless of the interaction device?*

2. *Can a virtual reality museum have an 'edutainment' experience pleasing enough to really use said museum?*

The answer to question one, if positive, would be grounds for further research in this direction to see if these results are not application specific and thus, if a design guideline could be deducted from it. In any case we have also measured the influence of interaction devices so another conclusion here would be which device might be best suited to use. Furthermore we had to take care of a possible drawback of the use of onscreen hints. They might have a negative influence on the immersion of a user since it is a visual reminder of the fact that the user is currently operating a computer program, this was to be addressed with the user as well.

The second question is harder to quantify. The problem of measuring fun and entertainment is extensively covered by (Wiberg, 2003) and (Jegers, 2003). By 'having an edutainment experience' it is meant if a virtual museum can have an added value on both entertainment and education over a normal museum. The interview, questionnaire and discussion will hopefully provide data for this.

7.2 HARDWARE SETUP

The test was done on the 'Podium' setup that is available at the Flexible Reality Learning centre. This choice to do it here instead of in the 'Cube' is based upon three arguments.

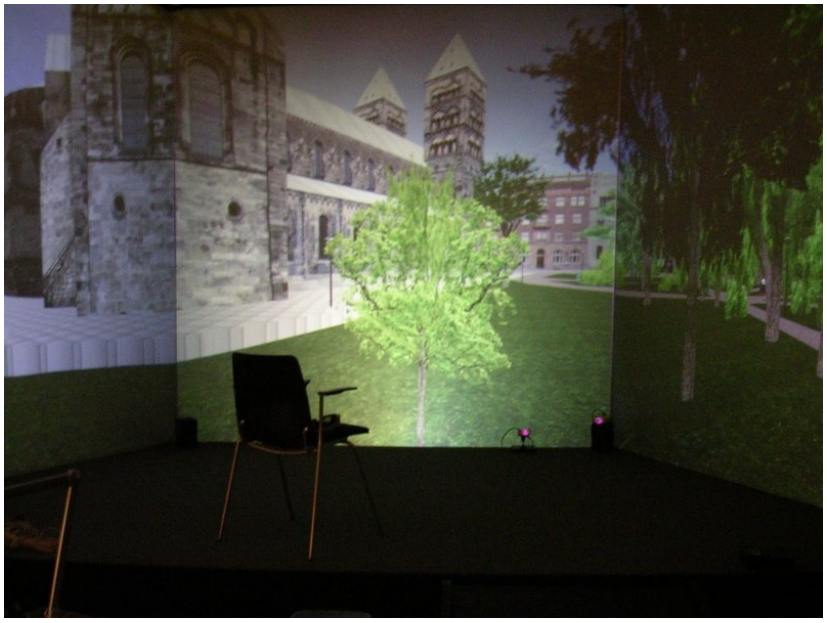

FIGURE 37 - THE VR PODIUM AT THE FLEXIBLE REALITY LAB

First of all, the odds of this application being used in conjunction with stereo projectors or a head mounted display are quite small because of the cost of the technology involved. A podium, on the other hand, still offers an immersive display where users can look around, but uses standard projection techniques such as beamers. One drawback is that the menu technique designed for this application was designed whilst keeping in mind that people usually reach for an object that is near them; this feeling of 'nearness' cannot be exploited on a podium setup and as such this menu in itself as an interaction technique might require more rigorous testing with a fully immersive display.

Secondly, it is technically more feasible. The cube is designed to be used with a head tracker and a motion sensing device. There is a template available which will always use these. Simply loading scenery in here is quite easily done. Adapting it for more specific interaction needs however is not always technically feasible or compatible with these devices. Furthermore it also remains to be seen if the ZoomBack technique still works, as it uses manipulation with one camera, and a cube simulation is split up into 4 different viewports of which the controls are quite rigorously implemented.

And finally, it is easier to observe subjects at the podium. There is a camera mounted on top that can register the subject that is using the podium and there is ample space behind the subject to set up a camera so you can also register what he or she is doing.

Despite all these arguments, it is still planned to adapt certain parts of the simulation, specifically the selection wheel, to the Cube so they could be tested later. However time did not allow it during the user study.

7.2.1 TEST REGISTRATION

As mentioned in these arguments the test is registered by using two cameras. Movements and feedback to those movements is explicit enough to warrant the use of only two cameras. One for the subject and one to film what is happening on the screen. These images are mixed together so it is easy to see what a subject is responding to. The camera on the subject is mainly meant to interpret the subject's facial expression and mood. A frustrated or elated expression can have serious implications for part of the design. The mixing is done using 'picture in picture' technique; the larger image is the user since his facial expressions and device actions are most important and the observer will easily be able to observe what's happening on the interface on a small screen.

Records and notes are kept of the answers to the questionnaire and the interview. The interviews are also filmed to make analysis of the answers more feasible, since every detail can be important for the improvement of the actual application. These notes are taken on a prepared answer sheet for the interviewer to make them more structured.

7.3 TEST GROUP

There is a variable that had to be tested 'between' subjects, which means that there are to be two disjunctive test groups. One group did the test without on screen hints, the other group was able to use on screen hints.

Because experience in 3D navigation and general 'tech savyness' can have an influence on the results both groups consisted out of a well mixed group of people with no overrepresentation of very experienced users in either one of the groups. This is generally influenced by gender, age and background (Wiberg, 2003) and as such the demographic makeup of both groups were about the same, and no large number of HCI researchers were overrepresented in a group.

There is still a lot of discussion on the ideal number of test subjects. It is argued by Nielsen that five persons are enough to reveal significant design flaws (Nielsen, 1993). However in response to this there are also claims that this does not cover complex application, to which VR applications certainly belong. In the end there is a general agreement that it is generally better to test a lot with a small group of users, than to expend all of your resources on a test with many subjects. Keeping this in mind and the fact that this is not a complete application by far, that the issue of on-screen hints is probably going to need more research and the resources available (manpower and time) are limited, the choice was made to keep the test groups relatively small for a statistical comparison; Each test group consisted of about eight persons. This is far more than the five suggested but since the goal is not only to reveal design flaws but also to make a quantitative comparison between different methods of interaction some more subjects were used.

7.4 TASK SET USED

The first part of the test is a short interview establishing the background of the user and their expectations of a virtual reality museum with respect to interaction, possibilities and media presented. These expectations were documented to compare if the end-results can live up to their expectations, and if not, what might be the cause of the difference. The background of users can be important here; it could have an influence in the statistical outcome of the test and the user's performance on the tasks. Age and gender were easy to determine, technical background involved some short questions about their experience in 3D Navigation, be it in games or as a

professional, their jobs and whether they own a Nintendo Wii and as such have experience with this form of interaction device.

The second part involved the user doing a set of tasks. During these tasks users were asked to 'think out loud' to gain insights as to why they are taking certain actions and where the design might suggest something other than expected by the designer. This also gave hints about problems users encounter and how one could correct them without leading the subject. These tasks will now be explained in greater detail.

7.4.1 TASK SET

To provide a user with a complete experience, a classification method by Pine II & Gilmore, as described in (Wiberg, 2003) is often used. It is called the 'four experiences' realm and was originally a framework developed to classify certain types of entertainment in the 'experience realm'. This experience realm was explained in chapter four. We will classify our tasks there to show that they cover most of it, and thus we are testing a design that offers a somewhat 'complete' experience. Tasks are based on the user tasks scenarios found in chapter five.

Task 1: Navigating whilst being able to look at and identify certain exhibits

This task will have the user explore the environment (in this case, the one exhibit room implemented in the application) and identify objects. To make the task more specific we ask the user to identify two certain objects, which are not immediately visible upon entering the room. Upon identifying the first object the user is to select it and unselect it again, after which he needs to select the second object. Once this is done the task is complete.

This task could be placed in the top half of the experience realm, providing some interaction for navigation and selection, and having the user absorb information rather than be immersed in it. When the user becomes more used to the interface this task will shift to the left and become more of an entertainment task.

Task 2: Finding content related to a selected exhibit

Here the user uses the selection wheel to finds two pieces of related content. These can be a movie, an audio file or an object that is related in some way. This task is done to give the user a feeling for the added value multimedia content can have to the simulation. Furthermore it is a test of the selection wheel that shows if it's natural to learn.

This task is about in the middle of the top half of the experience realm, it is both entertainment that is to be absorbed with a small choice by the user on the type of related content. The types of content currently implemented are all passive and as such no more mental effort is required.

Task 3: Finding a pop quiz about the 'sword' object and give an answer to the question

During this task the user builds upon what was learned in the previous tasks. The user starts in the start position and should move to the sword (this way learnability of the navigation interface is tested) and should select the sword to arrive at the quiz about the sword through the selection wheel. Using the information available through the interface (videos, text) the user should try and answer this question as accurately as possible. Since the user is now using two levels of the selection wheel we can check if the user understands this selection mechanism if it is applied consistently.

This task requires more mental effort of the user as well as use of the interface and could therefore be placed in the top right quadrant of the experience realm, being more of an educative experience than simply walking around and absorbing the environment.

Task 4: Entering the historical simulation and exploring it for a small while

The user should enter the historical simulation about the Drotten Church that is included in the application. This large virtual reality simulated can be accessed through the selection wheel. In this simulation a user can walk around a world that is modeled after the real Drotten Church including an outdoor environment.

This task focuses more on immersion and should appeal to the esthetic senses as well as a sense of experiencing something that cannot be experienced in the 'real world'. Therefore it is in the lower half. Depending on how enthusiastic people are about this one can choose to make this simulation more interactive.

The tasks as described were given in this order. There is a certain logic behind this. Next to providing the user with a full experience, these tasks also test the most important aspects of the application, as well as their learnability. The first task tests the user's ability to navigate the environment and selecting objects. This selection task is relatively easy. The second tasks builds upon the first task by having the user navigate to one of the found objects and then using the menu system to access two different types of content (thereby giving him a taste of the experience of having access to this kind of digital content). The third task then makes implicit use of this menu system through the question mechanism to see if the user understands the abstract concept of hierarchic selection wheels. The fourth task is purely for the experience of a simulation and the user's thoughts about this, thus will not be used for performance metrics, but might give design suggestions for the navigation.

If the user doesn't understand some of the mechanics involved in any task, he will have more difficulty doing the next task and thus we can see if the interface is easy to remember. Also, since every user did every task twice (once with each device) we can see if the interface itself is easy to use and learn, and if performance is improved after repetitions.

7.5 VARIABLES AND METRICS

The two variables that are important to this test are the interaction device used, and whether on-screen hints are used. These are the independent variables. They were compared in performance based on different metrics, the dependant variables. This data, amongst other things, has been used to answer the first question that is central to the usability test.

To measure efficiency of the best interaction method with regard to on screen hints one can use some traditional metrics on the tasks set. These metrics are: Time used to complete task, number of tasks successfully completed, number of wrong physical actions taken during the task and numbers of faults during the task. The difference between the latter two is that for the wrong physical action we mean that the user is using the interface device in a wrong way (for example: trying to use the motion sensitive controls on the nunchuk controller while in actuality the WiiMote should be used), while a 'fault' is the user doing something conceptually wrong (e.g. selecting 'audio' when he is looking for a video). In this way we can differentiate between errors and faults where the cause is a wrong understanding of the interface device, or a wrong understanding of the interface. These metrics are a collection that was considered relevant taken from (Stanney, 2002) and (Bowman et al., 2003), who in turn based their primary metrics on research done by Dix, et al. and Lampton, et al. More metrics were considered, for example the numbers of turns made in (Wallergård, 2007), but were found not to be relevant to a virtual reality museum.

The 'edutainment' experience is going to be somewhat harder to measure in exact terms. There is not much research available on this however, a PHD thesis written by C. Wiberg (Wiberg, 2003) also described in 4.6.2 handles almost exclusively on the definition and measurement of fun, in the context of websites. The introduction

however contains a very inspirational general part not specific to websites that might make it easier for us to classify what constitutes 'fun' in our application and how we can measure if it is actually fun and educational.

To see if the implemented version of our application is more educational than the simple passive absorption of information we have used a simple metric. As described in the task set all the users were asked to read something about an object. After that they were asked to read something about another object with the knowledge that they will have to do a quiz question after that. Then later, after the interview and discussion the user was presented with three questions. One is based on the first text, the second is based on the second text and the third uses information presented in the quiz-question the users answered. If a significant difference is measured in information retention we are able to conclude that active participation in a museum can lead to a more memorable experience. The users were asked to discuss on how 'educational' he thinks an experience like this might be; this will be discussed further during the design of the questionnaire.

To measure if our application is fun we have to, as mentioned before, have a better grasp of what exactly is the meaning of 'fun' in the context of the virtual museum application. Inspired by the answers of the initial user test and the first two chapters of Wiberg's thesis here are some of the most common attributes which can make an application 'fun' to use (though in some cases they enhance 'experience' rather than 'fun') in alphabetical order:

Attribute	Description
Aesthetic value	A pleasing graphical design can lead to a better experience (Also called Visceral Design).
Comfort	How physically comfortable a user is during the operation of the virtual museum.
Easy to learn	If the threshold is low users are more prone to pick up the application and get the most of it. Since this application is not a traditional game the concept of flow (Csíjszentmihályi, as quoted by Wiberg) is not applicable here, meaning that there is no point in making it harder to give people more of a challenge.
Immersion	People sometimes seek to escape thinking about everyday things; immersion can help in this aspect.
Interactivity/ Participation	Being able to influence the environment can provide a user with entertainment and a sense of influence.
Multimedia	Only having one type of media can lead to getting bored, while change is usually considered a very good thing in all walks of life.

TABLE 4: ELEMENTS OF 'FUN'

Now that we have a lower level definition of fun we can seek to measure each of these components.

As mentioned in chapter four not everything is as important to have 'fun'. We have defined four tasks and these four tasks should all cover certain sections of the above mentioned components. Easy to learn and comfort however should apply to all the tasks. Since most of these values are subjective we have tried to evaluate them using the questionnaire and the interview.

More subjective data will be gathered by simply observing the test subjects and asking them how they 'like' certain actions and abilities in the program. Also their responses to certain things designed to make the user feel good will be noted down and mentioned later in the results and conclusions. When users express that something is fun, challenging or in any other way interesting to them this should be noted to see if a pattern might perhaps emerge.

More objective data could be gathered using clinical metrics like heartbeat, an ECG or other satisfaction indicators produced by the human body. However this would involve quite a lot of equipment, time and manpower and was not within the scope of this project. However, if a final application is made, it could be interesting to compare these values with respect to a normal museum visit.

7.6 POST-TEST QUESTIONNAIRE AND INTERVIEW

To measure the more subjective metrics and to get inspiration for the improvement of the current design a post test questionnaire was presented to test subjects as well as an interview with prepared questions. Using this data we have then made some conclusions based upon quantitative data (the questionnaire) as well as qualitative data (the interview) on the subject of fun. The interview also provided us with insights and ideas on how to improve the application; in the best case the interviewed subjects had some very valuable ideas. However there are some matters we had to take into account when creating this questionnaire and interview. First we will treat some common issues in questionnaires:

First of all, as suggested on an online tutorial for questionnaire design[18] the best results are obtained when the subject feels completely at ease and anonymous. Therefore we must assure the subject that all answers will be treated confidentially. To this end his answer sheet will go into an anonymous envelope listing simply the test number so that it is still associable with the initial interview, but further being completely anonymous. This way we will hopefully avoid issues such as prestige bias where the user might be overly positive about his or her own experience.

Secondly, to gain quantifiable data we must make the questionnaire questions closed format. However we should still be careful with this. A disadvantage listed on the documentation of a course of questionnaire design[19] is that conclusions may be misleading because the subject might not be able to find the specific answer he is looking for. Therefore our questions were more specific, and the presented answers not vague. Where you could ask someone "How many historical simulations would you like to see" and present the subject with a scale varying from "Many" to "very little" it was preferred to be more specific; in this case that would mean having a better defined range of answers (e.g. "none", "1 or 2", "3 or 4", "5 or more").

The questions themselves must be non-leading, and where possible non hypothetical. While in some cases this cannot be avoided (for instance when asking whether a non-implemented feature would still add a lot to the current application) these cases should at least be very specific and well described, leaving not too much cognitive load on the subject taking the test and leading to reliable results.

For some aspects as defined in section 7.4 with regard to the 'fun' experience a user was presented with a scale on which to indicate his satisfaction with certain aspects of the simulation. Later on during the interview these

[18] http://www.cc.gatech.edu/classes/cs6751_97_winter/Topics/quest-design/ as visited in June, 2008

[19] http://www.tardis.ed.ac.uk/~kate/qmcweb/q6.htm as visited in June, 2008

answers and the reasoning behind it were talked about to discover what was good or bad about that particular aspect.

The interview is more open and therefore there is somewhat more freedom in the questions. However, as a rule, it is generally better to make no question which has 'yes' or 'no' as a valid answer. This way the user is forced to be specific about his answer and think about what they are going to say. There should also be room for discussion where the answers of a user can be challenged; however care should be taken not to offend the user or to present the challenge as a 'superior point of view'. Every answer is a valuable addition to the data and should be treated as such.

Also, during the interview the user's general opinion or feelings were gauged. This did not yield quantifiable results, but provided us with valuable insights as to why something does not work in the eyes of the user. Care should be taken with asking 'why' too much though. Wiberg describes a fictional session in a thesis where a technique is illustrated which is concerned with only asking 'why' at every answer. It is also described that this is usually quick to anger a user and should only be used in short sessions.

The final questionnaire and interview can be found in the appendices and will be important to understand the results presented in the next chapter.

8 RESULTS

In this chapter we will give an overview of the results obtained in the test described in the previous chapter. Conclusions about issues relevant to this thesis are made in the next chapter, partly based on the results that we are documenting here. We will explain how we addressed certain issues concerning reliability of the results and how we tried to keep the metrics 'pure'. After this we will show comparative results between the WiiMote and the SpaceBall, On Screen Hints and No On Screen Hints, and combinations of both. Furthermore the results of the questionnaire and the interview are presented.

8.1 TEST GROUP AND TEST ORDER

In the end we tested fifteen users for our application. Seven persons took the test without on screen hints and eight persons used on screen hints. We varied the order in which devices were used to make sure the learning effect had no influence on the results. Furthermore we also varied the order of the tests with on screen hints and without them. The relative inexperience of the observer could lead to 'better' tests in a later phase which we have tried to address in this manner.

The test group consisted mainly of students though some persons from the design faculty also participated as well as friends of people who participated earlier. This is reflected in the fact that the average age of the test group was 28.8 years old. The fact was stressed that a test-subject could not know anything about the application before they went into the test, thus leading to a rejection of some people who had a very colorful description by friends before applying for the test.

8.2 METRICS EXPLAINED

In the following paragraphs we will present tables with results; however first of all we would like to explain how certain metrics were measured during the actual analysis of the data gathered during the test described in chapter seven.

The time taken for a certain task proved to be a non uniform metric to measure. Certain persons would take a long time reading a piece of text, others would 'think out loud' very much and thus forgetting to actually use the device and sometimes something went wrong during the test were a reboot was necessary. Thus where we first assumed that timing could be automated, we had to manually redo those measurements from the videos. Only pure interaction time was measured. The time was stopped if the user stopped interacting (e.g. when he was reading a text, when he was explaining something in great detail) thus filtering out these effects that might change the results, to keep them as pure as possible. Furthermore if a user would take an extremely long time for completing a task simply by persevering in it, we would pick a reasonable maximum time to not skew our averages too much. By doing this consistently for every test subject very pure 'interaction time' is actually measured, thus providing us with a good performance metric. Time is always expressed in seconds.

Taking a look at the second metric, errors (pushing the wrong button for example), we noticed that a better definition was needed. For example when a user wants to move forward but moves down and tilts forward instead, are these two errors or only one? What we decided upon was the following: every time a user did something he obviously did not intend to do, we would mark that as one error. Hence if a user wants to go forward and he does something that does not make him go forward that is one error. If he repeats this (so he has stopped his wrong motion, but starts it again) this is counted as another error. Small variations such as a small nudge in the

wrong direction that the user doesn't even respond to were not counted, only if it was an obvious hindrance to the user. This worked quite well as users reacted quite strongly if something wrong happened.

The third metric, faults, was less hard. When persons did not 'get' what they were supposed to do in the interface this was quite easy to observe. By reminding them to think out loud users would already explain what they were looking for, and if a user is thinking in the wrong way about the interface an expert on the program, e.g. the test observer, can easily spot this and ask what is wrong. Using this we counted every time a user did not go on the correct course to attain his goal (e.g. looking for a virtual TV to open a movie instead of using the selection wheel) as a fault.

Using these methods of measurements we obtained a large number of results for our 15 test subjects which we will now present on a per-device basis as well as some different orderings which can show interesting results. The full results are available in Appendix D for readers who want to gain a more in-depth understanding of the results or perhaps find correlations not presented here.

8.3 Device results

For clarity's sake we present the result somewhat summarized, using comparative averages and standard deviations along with T-Tests to show the significance of the results. As the test groups were on the small side in case of the on screen hints the significance of these results are somewhat diminished.

8.3.1 WiiMote vs. SpaceBall

First of all we present the average (AVG) metrics and their standard deviation (SD) for all three tasked when one compares all the tests done with the SpaceBall (SB) vs. all the tests done with the WiiMote (WM).

	SB-AVG	SB-SD	WM-AVG	WM-SD
Task 1 – Time	116.87	65.73	78.80	45.50
Task 1 – Errors	4.47	3.58	3.07	2.55
Task 1 – Faults	0.53	0.83	0.67	0.98
Task 2 – Time	27.93	14.24	79.60	49.62
Task 2 – Errors	0.53	0.92	3.53	1.73
Task 2 – Faults	0.20	0.41	0.80	1.21
Task 3 – Time	37.13	12.38	46.00	22.84
Task 3 – Errors	0.40	0.74	0.33	1.05
Task 3 – Faults	0.33	0.62	0.47	0.64

TABLE 5 - WII MOTE VS. SPACEBALL METRICS RESULTS

The ρ-values for one and two tailed t-Tests for these comparisons:

SpaceBall vs. WiiMote	TT-2	TT-1
Task 1 - Time	0.02	0.01
Task 1 - Errors	0.11	0.06
Task 1 - Faults	0.72	0.36
Task 2 - Time	0.00	0.00
Task 2 - Errors	0.00	0.00
Task 2 - Faults	0.12	0.06

Task 3 - Time	0.23	0.12
Task 3 - Errors	0.85	0.42
Task 3 - Faults	0.63	0.32

TABLE 6 - WII MOTE VS. SPACEBALL T-TESTS

We present both one and two tailed t-Tests for the interested reader. The most important values are those of the two-tailed t-Tests as we had no expectations about the performance of any of the devices, hence we should have significant results that work 'both ways'. If a reader had his own expectations he can compare the significance of the result.

8.3.2 On Screen Hints vs. No On Screen Hints

Next we present the results obtained for On Screen Hints (OSH) compared to No On Screen Hints (NOSH) and their ρ-values.

	OSH-AVG	OSH-SD	NOSH-AVG	NOSH-SD
Task 1 - Time	83.63	51.89	114.07	63.86
Task 1 - Errors	3.00	2.48	4.64	3.65
Task 1 - Faults	0.44	0.73	0.79	1.05
Task 2 - Time	46.44	39.33	62.14	49.83
Task 2 - Errors	1.63	1.59	2.50	2.44
Task 2 - Faults	0.44	0.89	0.57	1.02
Task 3 - Time	40.81	15.32	42.43	22.36
Task 3 - Errors	0.25	0.58	0.50	1.16
Task 3 - Faults	0.44	0.73	0.36	0.50

TABLE 7 - ON SCREEN HINTS VS. NO ON SCREEN HINTS METRICS RESULTS

OSH vs. No OSH	TT-2	TT-1
Task 1 - Time	0.16	0.08
Task 1 - Errors	0.16	0.08
Task 1 - Faults	0.30	0.15
Task 2 - Time	0.34	0.17
Task 2 - Errors	0.25	0.12
Task 2 - Faults	0.70	0.35
Task 3 - Time	0.82	0.41
Task 3 - Errors	0.45	0.23
Task 3 - Faults	0.73	0.37

TABLE 8 - ON SCREEN HINTS VS. NO ON SCREEN HINTS T-TESTS

Since there is a very reasonable suspicion that users with On Screen Hints will perform better than those without them, the one tailed t-Tailed is the most important value here.

8.3.3 WiiMote vs. SpaceBall when using On Screen Hints

It can also be relevant to compare the influence of the on-screen hints on the difference between the devices hence we will present a comparison between the devices divided by the fact if they used on screen hints or not.

On Screen Hints	WM-AVG	WM-SD	SB-AVG	SB-SD
Task 1 - Time	71.63	48.86	95.63	55.25
Task 1 - Errors	2.63	2.97	3.38	2.00
Task 1 - Faults	0.63	0.92	0.25	0.46
Task 2 - Time	66.63	47.14	26.25	12.66
Task 2 - Errors	2.88	1.25	0.38	0.52
Task 2 - Faults	0.63	1.19	0.25	0.46
Task 3 - Time	45.88	18.39	35.75	10.31
Task 3 - Errors	0.13	0.35	0.38	0.74
Task 3 - Faults	0.50	0.76	0.38	0.74

TABLE 9 - WIIMOTE VS. SPACEBALL USING ON SCREEN HINTS METRICS RESULTS

OSH: WiiMote vs. SpaceBall	TT-2	TT-1
Task 1 - Time	0.08	0.04
Task 1 - Errors	0.43	0.22
Task 1 - Faults	0.28	0.14
Task 2 - Time	0.06	0.03
Task 2 - Errors	0.00	0.00
Task 2 - Faults	0.48	0.24
Task 3 - Time	0.23	0.12
Task 3 - Errors	0.35	0.18
Task 3 - Faults	0.78	0.39

TABLE 10 - WIIMOTE VS. SPACEBALL USING ON SCREEN HINTS T-TESTS

8.3.4 WiiMote vs. SpaceBall without On Screen Hints

NOSH	WM-AVG	WM-SD	SB-AVG	WB-SD
Task 1 - Time	87.00	43.56	141.14	72.32
Task 1 - Errors	3.57	2.07	5.71	4.68
Task 1 - Faults	0.71	1.11	0.86	1.07
Task 2 - Time	94.43	51.67	29.86	16.67
Task 2 - Errors	4.29	1.98	0.71	1.25
Task 2 - Faults	1.00	1.29	0.14	0.38
Task 3 - Time	48.67	30.56	38.71	15.11
Task 3 - Errors	0.67	1.63	0.43	0.79
Task 3 - Faults	0.50	0.55	0.29	0.49

TABLE 11 - WIIMOTE VS. SPACEBALL WITHOUT ON SCREEN HINTS METRICS RESULTS

No OSH: WiiMote vs. SpaceBall	TT-2	TT-1
Task 1 - Time	0.10	0.05
Task 1 - Errors	0.20	0.10

Task 1 - Faults	0.85	0.42
Task 2 - Time	0.02	0.01
Task 2 - Errors	0.01	0.00
Task 2 - Faults	0.17	0.09
Task 3 - Time	0.59	0.30
Task 3 - Errors	0.85	0.42
Task 3 - Faults	0.69	0.34

TABLE 12 - WIIMOTE VS. SPACEBALL WITHOUT ON SCREEN HINTS T-TESTS

8.3.5 First round of tasks vs. Second round of tasks

Every user did the task set two times and as such we should also compare the averages for the first round compared to the second round to see if learnability of the interface is high, especially considering the metrics time and faults.

Round 1 vs. Round 2	AVG Round 1	SD Round 1	AVG Round 2	SD Round 2
Task 1 - Time	110.27	60.47	85.40	56.37
Task 1 - Errors	4.13	2.83	3.40	3.48
Task 1 - Faults	1.00	1.07	0.20	0.41
Task 2 - Time	70.80	54.28	36.73	22.93
Task 2 - Errors	2.53	2.33	1.53	1.64
Task 2 - Faults	1.00	1.13	0.00	0.00
Task 3 - Time	50.53	20.13	32.60	11.81
Task 3 - Errors	0.53	1.13	0.20	0.56
Task 3 - Faults	0.80	0.68	0.00	0.00

TABLE 13 - ROUND 1 OF TASKS VS. ROUND 2 OF TASKS METRICS RESULTS

Round 1 vs. Round 2	TT-2	TT-1
Task 1 - Time	0.15	0.07
Task 1 - Errors	0.42	0.21
Task 1 - Faults	0.02	0.01
Task 2 - Time	0.07	0.03
Task 2 - Errors	0.29	0.14
Task 2 - Faults	0.00	0.00
Task 3 - Time	0.01	0.00
Task 3 - Errors	0.33	0.17
Task 3 - Faults	0.00	0.00

TABLE 14 - ROUND 1 OF TASKS VS. ROUND 2 OF TASKS T-TESTS

8.4 QUESTIONNAIRE RESULTS

Here we will present an overview of the answers to the questionnaire, or more specifically the averages. Presenting all the answers would be quite a lot of data hence we will suffice with averages which are also used for conclusions. The results are presented mostly with numbers. Whenever the user had the choice between two options (a yes or a no, SpaceBall or WiiMote, etc.) we will mention the option that was named the most as well as the number of times it was crossed, for all 15 tests performed. For the one-to-five scales we will use the numbers one to five corresponding left-to-right with the options on the questionnaire. In almost every case higher numbers are better. An exception is the 'order of difficulty'. Here users were asked to order tasks in order of difficulty, using the highest number for the hardest task. Hence the higher this score is, the harder a user found the task.

Questionnaire results	
Aesthetic	
Would you have liked a bigger room?	8 times no
How much bigger?	2.29
Would you have liked more rooms?	13 times yes
How Many?	2.83
Was the LOD Adequate?	1.73
Comfort	
Which would you prefer for the SpaceBall: Sit/Stand	12.5 sit
Which would you prefer for the WiiMote: Sit/Stand	10 stand
Did you notice any strain(SB)	12 no
Did you notice any strain(WM)	14 no
If there was strain-More than using a computer?	4.00
Learning	
Which was the easier device to understand	8 WiiMote
Order of Difficulty	
SB-Walking Around	3.53
SB-Looking Around	2.47
SB-Operating wheel	2.33
SB-Selecting	1.67
WM-Walking Around	1.87

WM-Looking Around	3.00
WM-Operating wheel	3.33
WM-Selecting	1.80
How easy was it to:	
Learn the Educational Game	4.07
Enter the Historical Simulation	4.27
Find related content	4.47
Stroll around	3.87
Interactivity	
Did you miss the option to manipulate objects?	8 times no
How natural was the implementation of the SpaceBall?	3.07
How natural was the implementation of the WiiMote?	3.17
Immersion	
Did this feel like a real museum visit?	2.43
Was the Historical Simulation believable?	12 times yes
Did you find the immersion important when:	
Walking through the museum?	12 times yes
Exploring information about exhibits?	8 times no
In the Historical Simulation	14 times yes
Did the On Screen Hints influence your sense of immersion?	7 (out of 8 OSH-tests) times no
If yes, how much?	3.00
Appreciation of Features	
Browsing related content	4.53
Small quiz questions	3.93
The Historical Simulation	4.20
The ZoomBack navigation technique	4.73

TABLE 15 - QUESTIONNAIRE RESULTS

8.5 INTERVIEW RESULTS

We will try and give a short overview of the most important answers given during the interview. Of course a complete transcript would be too much hence we will try and give the common denominators among the answers, like we did for the initial user study performed at the start of this project.

What was the thing that appealed to you most in this simulation you just did?

Answers to this question differed quite a lot; hence it is difficult to find some common ground. What was never mentioned however could be interesting. Users never seemed to think about the ZoomBack technique in this open question format. Another interesting result was that a lot of user just liked using the SpaceBall or the WiiMote in any context, since they were 'new devices' to them so they were mentioned a couple of times. If any common denominator has to be appointed here it would be the historical simulation and the freedom enjoyed in the simulation. Mostly though the users mentioned that a lot of what they did was enjoyable.

And what appealed to you the least?

Here the answer given was very often the same, namely troubles with orienting or navigating the viewpoint leading to disorientation and embarrassment. This often corresponded to the task the users found the hardest in their particular tests. Never mentioned were the possibilities of the museum. A special mention goes to one user who found it annoying that a virtual museum was so empty, while it could be designed as 'so much more'.

How did you like the two devices?

Often this question corresponded much to the questionnaire answers, the users used their answers there to lead their reasoning here. In the end the WiiMote was often mentioned as a very fun device, with the gestures as the biggest plus, even though they were also the biggest problem. The SpaceBall was often surprisingly good to people, especially since they expected it to be a difficult device due to the bulky look and their unfamiliarity with this device. The freelook mode implemented on the WiiMote was also mentioned often as being way too sensitive. Some comparisons here were often made on which device they preferred, sometimes not corresponding to the device that was easiest to learn in the questionnaire. This was due to the fact that with training persons started to appreciate the other device more. This leaned in both the direction of the WiiMote and SpaceBall though.

What did you think of the educational game?

Most people seemed to think it was quite good. The fact that you could choose to do it instead of being forced to do it was very much appreciated. What was interesting to see was how men appreciated it more to keep track of scores and compare them, thus more or less confirming the fact that men might be more competitive. Women however sometimes would like it as well, but at the least did not mind this. It was mentioned sometimes that a somewhat more intricate game would be liked. One person really did not like the game, this could be due to the fact that this person didn't get a single question right implying that questions should not be too hard.

What about the big simulation?

The consensus here was that it was very much fun, and the atmosphere created there was very much appreciated (through fog on a graveyard around a church) but it could definitely use some more design work and mostly some simple interactions. Users often wanted to be able to get to know more about certain objects that stood out (a tapestry, a book on an altar, the church bells, etc.). Sound would also be very welcome.

How would you compare this experience to a 'true immersive experience' (e.g. with stereoscopic glasses)?

This was a hypothetical question and as such was often difficult to answer. We tried asking for more specific aspects such as immersion, or the selection wheel. The selection wheel using stereoscopic display was welcomed with enthusiasm and was described as "probably very cool". In the end most people had a hard time imagining it and didn't see much added value for the normal museum. The immersive experience in the historical simulation would be much improved however.

Is there anything regarding the interactiveness which you would like to see extended?

Manipulation was often mentioned here, though that might have something to do with the fact we asked users to consider this in the questionnaire. Others commented they would just like 'more'. More specifically people would like to pick up things that afforded picking up such as the hammer and the sword. A rather nice answer was interactive diagrams/pictures which you could browse through or get some more information from in an interactive way.

What about new possibilities in this museum?

This is quite a big question and as such users sometimes had no idea what to answer. However some very nice ideas were presented. One user would like to see a space station modeled like this and then have the information retrieval options offered by this framework. Another suggested traveling through the human body. All these suggestions were basically specific versions of making the 'exhibition space' more context sensitive. Furthermore users would like background sounds when they selected something. Another suggestion was personalized exhibitions which ties in to some research done in chapter four.

How did this experience measure up to a real museum visit in regards to education and fun?

Often people responded this was more fun, due to extended possibilities. A comment was made however was that this didn't create the feeling of a connection with history so apparent in a normal museum since the exhibits there are actually 'real'. As for educational, people would argue this could be more educational due to the possibility for information retrieval, however one had to create a sense of 'wholeness' and once again there was the point that a lack of 'connection' with history might make it harder to care about everything. Very much liked with regards to a normal museum was the fact you could make choices about what sort of info to look at per exhibit, thus creating a personalized experience.

And finally, would you actually use an application like this if you came across it?

Only two persons thought they wouldn't. All the others answered with a definite yes (though some made the observation that they would only do it if the exhibition suited their interests, which seems logical). The persons who did not like it had two reasons. One said she just didn't much like using computers to go to a museum and probably never would (though she really liked the technology involved, just not in a museum context). The other said she would feel a bit embarrassed if others could see her try out things in this computer program and that threshold was too high.

9 CONCLUSION, DISCUSSION AND FURTHER WORK

Now that the results are known we can draw some conclusions from these results. We will start by comparing the performance of the two interface devices used, after which we will discuss how the on screen hints influence this performance. This should answer the first question whether on screen hints significantly improve user performance regardless of interaction device.

After that we will discuss the results and observations about the interface. Note here that we mean the interface as in the way the program is designed to work, not the mapping of a device to the interface. During the analysis of the video many observations were made about the actual interface, which are further explored in the questionnaire. We will show here what did and what did not work in our interface and how this supports our interface being 'fun' and 'educational'. This will be followed up by a section which lists improvements that can be done to address these shortcomings.

After that we will deal with conclusions that were made during this project that were not specifically stated as goals, but that were stumbled upon during the entire project.

We will end this chapter and thesis by giving some suggestions for further work that can still be done and in some cases still needs to be done to confirm certain suspicions, as well as future projects that can be done as follow-ups on this project.

9.1 INTERACTION DEVICES: SPACEBALL VS. WIIMOTE

First of all I will start of by showing the conclusions I have been able to draw about how the WiiMote and the SpaceBall compared to each other based on metrics and how the interface influenced this. Afterwards some more general observations that were made during the analysis of the user study videos will be mentioned for completeness.

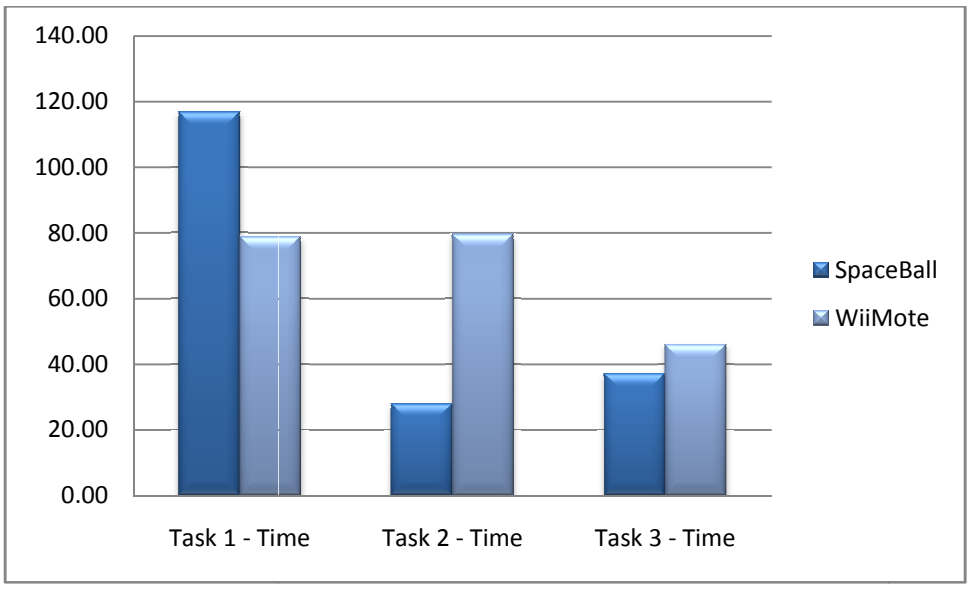

FIGURE 38 - GRAPH SHOWING PERFORMANCE TIMES FOR THE SPACEBALL VS. THE WIIMOTE IN SECONDS

The first conclusion that we are able to arrive at is that the WiiMote is better for navigation than the SpaceBall. This conclusion is based upon the significant difference in the average time needed to complete task one and the amounts of errors made during this task, the WiiMote scores better on both counts. Task one was almost purely about navigation (finding two objects in a new environment and selecting them) and therefore it is safe to assume that the WiiMote worked much better for navigation. This is probably due to the fact that we used a joystick to navigate, which gives a user 2 DOF which is easier to control than the 6 DOF a user is controlling using the SpaceBall; however the possibilities with the WiiMote for orientation and movement are still the same as those with the SpaceBall. They are just better divided across the interface device by having an (isometric) joystick for movement input and another (isomorphic) part for orientation.

For menu selection using a selection wheel however, the situation is reversed. Here the SpaceBall scores significantly better on the count of both errors and time taken. The main cause of this was the unfamiliarity with the gesture interface of the WiiMote and the fact that there was no alternative to this gesture system, hence people would assume the selection wheel didn't work since it didn't respond to their actions. However even the alternative actions taken were not as uniform as the SpaceBall; the SpaceBall seemed to have a very natural way of turning the wheel for people. This is also reflected in the questionnaire where operating the selection wheel is listed as the easiest task to do with the SpaceBall, while with the WiiMote this is the hardest task to do.

The third task which was included to test learnability of both the navigation and the menu system (though it focused most on the menu) differs a lot less. The significance of these differences is not very great according to a t-Test performed on the results ($\rho = 0.23$ for a two-tailed test). We can therefore conclude that both devices are comfortable to use after training without any significant difference between the two of them. It must be noted however that the WiiMote was experienced as more fun due to the gesture interface which "feels powerful" according to one user whose opinion was mirrored by many other different users.

Finally the real question: which interaction device would be recommended? In this case I would argue in favour of the WiiMote, though a SpaceBall is by no means a bad choice, and if support or compatibility is an issue these are good practical arguments. The WiiMote however is the cheaper solution, and if one extends the interaction for navigating menu instead of using only gestures for it (thus mapping multiple device actions to one interface action) one retains the possibility of later extensions using the gesture interface that can make a museum much more fun and immersive, while having a good basis. When people finally found out about the gestures this was usually paired with some elation, indicating that the fun-factor is certainly present for the WiiMote which tips the scale in its favour.

These results are taken in the Virtual Museum Framework but as the tasks described are quite general it leads us to believe that they might hold true for other applications where navigation and menu selection are important tasks as well. While we cannot conclude this conclusively there is certainly some basis to assume that the devices will work in the same way.

9.2 On Screen Hints and their influence

A question I set out to answer was whether on screen hints influenced performance regardless of device. We have found out this is the case, though it influences tasks that are already known to the user and in which the user is trained a lot less and only helps with 'problematic' tasks. We will now elaborate on this a bit more.

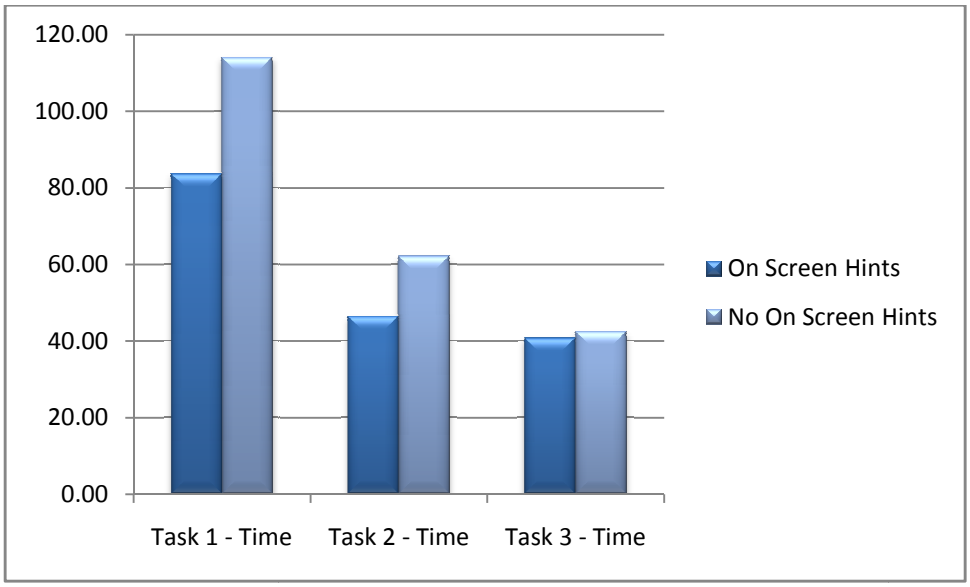

FIGURE 39 - GRAPH SHOWING PERFORMANCE TIMES FOR ON SCREEN HINTS VS. NO ON SCREEN HINTS IN SECONDS

A basic comparison between tests performed using on screen hints and tests performed without on screen hints shows a big difference in both time and errors for task one and task two. Task three however, where the user is already trained does not show this difference (as a matter of fact, there is almost no difference). A problem here is that the number of trials is more limited since we have split up the test groups and were thus left with a much smaller group. Therefore the results are somewhat less significant as is reflected in the slightly higher t-Test results. While one could very reasonably expect that on screen hints improve performance, thus using the one tailed t-Test, this still only gives us $\rho = 0.08$. This is quite good for such a small test group however and we can at least conclude there is a big indication that on-screen hints do their work.

When one takes a look at the results per device the number of trials becomes even smaller, however one can still observe that in the case of the WiiMote, the on screen hints improve task two quite a lot, which was the weakness of the WiiMote, whereas they don't really have a significant result on task one. In the case of the SpaceBall this is reversed, task one is much better off using on screen hints and task two is not significantly improved. Furthermore the results are quite the same per device as those for both devices, showing that on-screen hints do help, regardless of the device.

A more general conclusion would than be, that on screen hints improve performance in the case that a task is already inherently difficult, regardless of the interaction device. Of course one should still aim to design one's interface as natural as possible, but when this is simply hard to do, or when training a user, on screen hints can certainly be a good tool to assist the user. A corollary would be that one could perhaps use on screen hints as a measurement tool to gain insight how natural an interface is to grasp. If the on screen hints do not significantly improve user performance, assuming the hints are well designed, the task is apparently very natural (or so inherently difficult that even hints don't help).

9.3 VR Museum Framework design

An overall conclusion I can make about the interface is that performance wise it seems to be very well geared towards the tasks that are asked of the user. This is concluded from the very low amount of faults made by users. In the final task there were 0 faults and 3 errors made by 15 users across the board. As can be observed from the t-Tests comparing the first round of tests to the second tests the number faults is always very significantly reduced (observing $\rho = 0.00$ multiple times). This leads me to conclude that most of the interface is very learnable and tasks are easily repeated. This is also reflected in the fact that when the tasks are repeated a second time there are virtually no faults made (3 faults in 45 trials). The interface as it is currently designed seems to work very well, which is also supported by the questionnaire, where the appreciation for all tasks was quite high, and the ease of learning was also ranked very satisfactory.

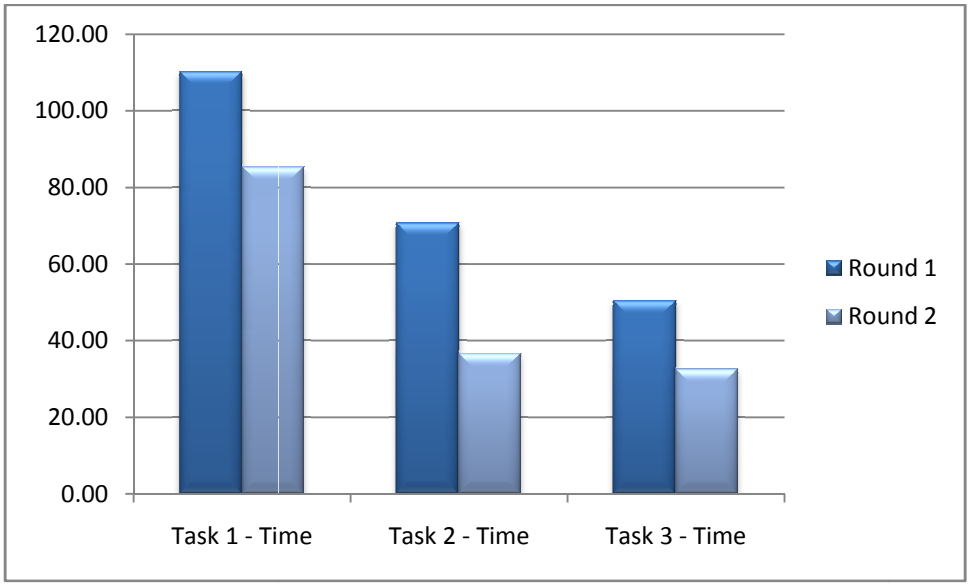

FIGURE 40 – GRAPH SHOWING PERFORMANCE TIMES FOR THE FIRST TIME TASKS WERE DONE VS. THE SECOND TIME THEY WERE DONE IN SECONDS

Of course this leaves us with some questions. Why is it working so well? What can still be improved? Besides the fact that it is efficient, is it also what the users actually want and appreciate? To answer these questions and give some grounds for discussion we will take a closer look at what we observed during the user trials and the answers given to the questionnaire and interviews and present a line of argumentation as to why the end result worked quite well.

9.3.1 Expectations vs. Possibilities

First off we will take a look at what people expect when you introduce an abstract concept such as a 'virtual museum' to them. Summarized from the results we can conclude that people expect to have an environment that they can navigate themselves, rich in highly detailed content with simple games such as a quiz or a puzzle. Preferably there are other virtual avatars walking around that can be interacted with and occasionally there are sounds. A tour would be appreciated.

If we look at what is possible in our framework we can come to the conclusion that the basics are there. The interface allows a user to navigate by himself, supports having highly detailed content in different contexts and contains mechanics to include quiz questions. It is however lacking in any framework for sounds, and no avatars

are included nor is there a framework provided for them. Many users however showed an understanding or even a hesitance when mentioning complex systems such as avatar and classified it as something they would want but not expect. Since we implemented what users would expect this means they will not be disappointed when using the museum, which adds to their experience.

9.3.2 What worked

Taking a look at what worked we can once again look at navigation and selection. Moving around was very easily done and it was much appreciated that the user could do this themselves. Because of the use of the ZoomBack technique the navigation didn't have to be very precise to actually reach intended goals in the application. This fact is also shown by the high appreciation for this feature in the questionnaire (it is the most appreciated feature in the application).

Selection of an object sometimes posed problems and we will come back to that in the next paragraph; however menu selection proved to be a very simple and effective concept. The selection wheel worked as intended and in only two trials did users not grasp that almost everything in the application was accessible through the selection wheel. In the end users seemed to expect it when entering the simulation, since the interface element was so consistently applied. The fact that users found it easy to understand how to access tasks available in the selection wheel is shown in the questionnaire results where the understanding of tasks relevant to the selection wheel score above a four out of five, whereas the only task to score just barely under a four (and this is more correlated to users who use the SpaceBall) is strolling around. Furthermore the fact that 'browsing related content' was rated on a second place of the features of this particular virtual museum framework leads us to conclude that users are happy that it's there, and also about the way it is implemented.

Besides basic VR tasks such as navigation and selection there were also features that focused more on the experience, and these seemed to have the desired effect. The little quiz that was included increased information retention as shown by the simple test that we did. All the users answered the quiz question correctly when it was asked later (a user that I spoke to a month later still remembered the answer) while questions about information in the museum that was not quizzed by the application were much harder and often wrongly answered. A preliminary conclusion would be that quizzes do increase information retention, but our testing on this was quite limited. The quiz was the least appreciated but with a score of almost four out of five it still scored quite highly showing it is not only educational, but also fun. Also users really cared about getting the questions right and were clearly very happy when the 'You are correct!' screen appeared.

The historical simulation was no longer the most appreciated task as it was during the initial user study but the fact that it was there and had quite an atmosphere using fog on a graveyard was very much appreciated. Exploring the simulation and flying through it were always met with happy attitudes. As we observed the users we can definitely conclude that this has the biggest 'fun-factor' in the simulation and if any design efforts are to be focused in creating an attractive exhibition, this would be an important part of it.

9.3.3 What didn't work

Unfortunately not everything went smoothly. Whilst navigation was designed pretty much as user wanted it, thus manipulating the viewpoint themselves, and through use of the ZoomBack technique, users still expected a tour or a talk sometimes to aid them in their navigation and choices of devices. We can conclude that our mode of navigation is enough, but can certainly be extended using these elements.

Selection had a problem. Our application was designed to select objects that were in the center of the screen but it seems users do not always grasp this. At first it seems obvious but later it was translated to 'looking at it', hence if

it would be on the side, but the user would be looking at the object and be close enough to it, he still expected the object to be selected when pressing select. It seems that the chosen selection method is not optimal.

The reason the historical simulation was no longer the most appreciated feature seems to be a discrepancy between expectations and what was offered. Often users would expect interaction to be possible in the historical simulation similar to that offered in the virtual museum room. It is our conclusion that the historical simulation should definitely contain interactive elements and sounds, perhaps even fully replacing the museum room to combine the best of both worlds. Of course one could also extend this to avatars or animations but users seemed quite understanding that these were lacking hence it is probably not necessary to include them to significantly improve user experience.

9.4 POSSIBLE IMPROVEMENTS

In this section we will list some improvements that seem to address issues that arose during the observation and analysis of the results. This of course is not definitive or conclusive but can be used as guidelines for improving the framework as it is right now during later projects. The inclusion of such a list was also recommended by Tromp et al. as described in chapter four (Tromp, Steed, & Wilson, 2003). The list of improvements contains no important results as to how the devices and interface operated. These questions were answered in the previous section.

9.4.1 INTERFACE SUGGESTIONS

9.4.1.1 NAVIGATION

While navigation was fairly comfortable there is a lot of room for improvement here on the subject of degrees of freedom. Even though we did not allow for the rolling movement on the WiiMote (it seemed weird to leave it out on the SpaceBall though) users would still end up in a rolled position, because they could look up, and then look sideways. Since the 'navigation plane' would still be pointing up, rotating sideways would in effect make you end up in a rolled position. To address this issue I would suggest constraining the navigation plane to a horizontal position, while still allowing the camera to orientate itself freely (thus one would not update the pitch and the roll of the navigation plane as defined in chapter five). In effect this means that you can no longer walk up diagonally or move sideways diagonally. It is however still possible to attain any position and orientation as one can still move upwards and downwards, left and right, backwards and forwards and orientate the camera in any way. In effect one could compare the navigation plane to the human body, which does the movement, while the orientation of the camera can be seen as the head mounted on the body looking around.

In any case, users can always end up in confusing positions hence it would make a lot of sense to add a reset action to the interface which moves and orientates the viewpoint back to ground level without any roll or pitch. Another big help that would prevent users from going too far 'out of bounds' would be to enable collision detection. This would require either fixing the bug which prevented us from implementing this, or using another implementation framework.

To improve selection of devices one could give a better visual cue when they are selectable instead of the hint that is currently showed on top of the screen. At the moment the system uses 'intersection' cubes to lower the computational load of a user looking at an object (there is a transparent cube around exhibits, if the user's gaze intersects this cube it is defined as 'looking at the object'). One could raise the opacity of these cubes a bit to show that the object is currently selectable.

9.4.1.2 Overall Design

The overall design of the application can certainly use work. The level of detail was often experienced as much too low and the exhibits certainly need improvement. However next to the obvious adding of more polygons there were also some rather nice suggestions towards screens. The information screens could also be a lot more stylized to reflect the theme of the exhibition, thus creating a more 'whole' sense in a finished application. Also by putting the objects in their actual historical environment one combines the historical simulation with the advantages of a virtual museum room. This would mean 'losing' a general exhibition area, but the resulting exhibition would be quite well tailored to the theme of the virtual exhibition.

9.4.1.3 Improving the selection wheel

There are several shortcomings still in the selection wheel even though it worked quite well, sometimes there were still misunderstandings.

To make the goal of the selection wheel more obvious when the selection wheel pops up, one could add arrows to the left and the right, to give a more visual clue that one can actually scroll left and right using this wheel. Furthermore if a selection wheel is opened on top of another wheel one might give it a different color to show more clearly that there is a new selection wheel which replaced the old one. The animations that are provided in the application as it is now were sometimes not enough. Another more aesthetic selection was to not have the abstract spheres used now in the selection wheel, but to also create virtual icons for what they represent (e.g. the 'audio' option could be represented by a speaker). This would certainly create a nice sense of immersion and also gives more clues as to the functionality of the options you are looking at without actually going there in the selection wheel.

9.4.1.4 The information screens

It was not always clear that a user could and should go back from an information screen before doing anything else. This can be addressed by adding a hint inside the information screen, instead of the more general 'Press B to go back' which is present now at the bottom of the screen. Furthermore instead of just text on the left side the information screen could also present several images on the right side (perhaps even a slideshow) to visualize certain parts of the text without being intrusive.

9.4.1.5 The Quiz/The Game

At the moment the quiz seems sufficient for users, however some more features can be added. A score could be tracked and also coupled to a certain 'grade'. This might be especially fun for kids; if for example, they get 60% of the questions right they could get a little diploma, thus stimulating them to learn and answer correctly. The system could make some more directed suggestions on how the score could be improved or where an answer could be found instead of the general answer it gives now. Furthermore, when designing these questions one should take care that they are actually answerable using the text and that they are non-ambiguous (one of the questions in the test was a bit ambiguous which lead to some irritation for some users).

Another suggestion would be to add more puzzle-like games such as putting object in the right order, image puzzles or anything else you can think of. The framework would need much extending though to allow these interactions to take place.

9.4.2 DEVICE SUGGESTIONS

9.4.2.1 WIIMOTE IMPROVEMENTS

The WiiMote can use a number of improvements, mainly in the area of orientation of the viewpoint and operating the selection wheel.

Orientating the viewpoint is too sensitive at the moment. It is now done by pressing A+B and then moving the WiiMote. If it is tilted up the viewpoint will tilt up and if it is tilted down it will tilt down. The sensitivity of these actions should be lower. Furthermore the baseline for the 'tilt' (the zero-tilt so to speak) should be the tilt which the WiiMote is at when a user presses A+B, so that it is relative to that initial tilt, instead of relative to holding the WiiMote parallel to the ground.

Menu operation at this time can only be done using gestures. While this is fun it should also be possible using the directional pad on the WiiMote or by moving the nunchuk joystick left or right when in 'selected mode'. This also applies to the select and unselect buttons. At this time these actions are only done by the A and the B button, however the buttons on the Nunchuk should also map to these actions, as they were often used during the tests without on screen hints.

To make the wheel turn more comfortably and use the analogue nature of a gesture one could also differentiate between a big swing and a little swing on the wheel. A big swing would than skip some options on the wheel where as a little swing would just go to the next option. Of course it is important that the system can differentiate between these two gestures quite accurately.

9.4.2.2 SPACEBALL IMPROVEMENTS

As mentioned navigation was a big problem for the SpaceBall. Improvements here are somewhat harder since this device inherently has 6 DOF to control and it is hard to limit it on the device without limiting the interface to 2DOF. However, changing the interface to not use 6DOF but '4 DOF' should already prevent many of the observed problems. Another thing that needs a closer look is the thresholds at which a user tilts forwards. This happened very often when the user was just moving forward. Unfortunately the action for moving forward was reasonably well split between tilting and moving the SpaceBall forward, and since in this application it is important that pitching the viewpoint down and up is possible, it seems not possible to map both actions on this device to moving forward without losing an intuitive way of tilting the viewpoint.

A suggestion made by a user was to use the SpaceBall sideways. In this way it is much easier and ergonomic to use two hands to operate the SpaceBall, or rather this version of it (the Trioc 3D), making the forward motion much easier and relieving some of the strain that some users experienced. Of course all the controls would have to be flipped 90 degrees in this case.

9.5 RELATED RESULTS

Besides the results obtained for answering the goals we had when we set out to do this project, some other conclusions can be made about the process used to attain the results, and the user study also showed some promising results for further research that are not directly related to the questions we set out to answer.

In chapter five it was mentioned that we did not use UML models for explaining the user task scenarios. While one might think a use case model here would be sufficient we noticed this was not the case. We concluded that the functional approach of these use case models did not allow for enough concerns about user experience to be documented. We solved this using three elements to explain what was important for any use case namely the user

tasks, the system responsibility for these tasks and 'important issues' for the experience. We do not mean to suggest that this is necessarily the best way, but there is a gap here that could certainly use more research.

We used a design process described in chapter four by Gabbard, Hix and Swan as mentioned in (Bowman, Kruijff, LaViola, & Poupyrev, 2004). However using other suggestions made by Bowman (Bowman, Gabbard, & Hix, 2001) we extended the final usability testing done in this design process with extensive formative and summative testing supported by less formal methods such as questionnaires and interviews. This gave us a very complete picture of what was wrong with our framework and how users might want to see it improved, giving us the ability to include some solid redesign suggestions in this chapter, which was mentioned as being necessary by Tromp et al. in (Tromp, Steed, & Wilson, 2003).

Furthermore during the evaluation phase, where we tried to assess how 'fun' our application was, we had some moderate success. However nothing conclusive could really be said about our results. We noticed that measuring fun is just inherently very difficult due to the many factors involved that could influence this, down to the mood a test subject is in on any particular day. This fact also makes it hard to attain reliable results using a low number of test persons, even if one uses clinical metrics.

Much research so far has been done in comparing 6 DOF navigation vs. 2 DOF navigation, which are both traditional methods in 3D. However with the rising popularity of computer games a 3 DOF mode of navigation has become increasingly popular (where a user would be able to pitch the viewpoint up and down). During our tests we included that if one extends this with being able to move up and down as well, in effect using 4 DOF, users are still very able to conceptually grasp what is going on. Furthermore it is possible to attain any possible viewpoint as long as one does not include roll, which was not very much appreciated anyway. Hence we can conclude that 4 DOF navigation has a lot of promise for application where freedom of movement is important while not making navigation too difficult for non-expert users.

Another conclusion that we have been able to make about the navigation is that using camera-based steering in an immersive projection display can lead to some confusion. Test subjects were often observed looking to the sides and then expecting the camera to move in the direction in which they were looking, while at other times they would expect it to just move forward (for example they would be looking at the left screen and would expect the forward direction to be 'to the left', while on other moments they would (correctly) assume that the forward movement went in the direction of the center screen).

9.6 Further Work

An interface element that worked very well during our user study was the developed selection wheel. This selection technique can certainly use some more work and is deserving of some research in this direction. One can think here about extending it for more complex structures, for example to see if one could visualize 2-dimensional data or even complex ZigZag structures which would create a very usable solution for complex data visualization. The usability attributes can also use work, such as adding virtual icons or defining which placement of the selection wheel would be best (for example, one could also lay the selection wheel out around someone and use pointing to select options, which might increase selection speed). Furthermore one could see if it works as a general selection menu, for which a testbed evaluation would have to be done, comparing it to existing techniques.

This framework could be used to passively build a user profile of anyone using it and as such allowing for the creation of a completely personalized experience. Research would be needed on how to passively gather certain indicators about a person's preferences and how one could use this data to create a pleasing experience for the user.

The educational game as it is, is very simple. More research would be needed on how to integrate small games in simulations like this. It has been shown that they can be very educational and to greatly support learning tasks. However there is currently a lack of examples or ideas for such small games and if one were to invest time into this they could be a very welcome addition to this framework, also adding to the much wished for variety.

As mentioned 4 DOF seems to be a very promising mode of navigation. More comparative studies focused on navigation tasks comparing this to the traditional 2 DOF and 6 DOF navigation would be needed to see if it might be better for certain applications that may require accurate positioning of the viewpoint using intuitive techniques.

The On Screen hints have been shown to be promising in that they enhance performance of difficult task. A small evaluation has been done on their influence on immersion but this was far from conclusive. There are a lot of variables that can be varied such as position, size, color, etc. Furthermore instead of pictures one could use on screen movies to explain gestures. Another possibility would be to add 3D objects that give hints about the use of onscreen devices. A comparison between these aspects is needed to arrive at some definitive conclusions about the use of on screen hints in any general application.

The user study here focused on individual users. However an advantage to visiting museums is the fact that you can do it with friends. Further research would be needed on how to extend this framework to better support collaborative users. In the future one could even think about networked applications, however for now it seems better to focus efforts on having multiple users work together/explore together in a virtual museum using projection displays.

10 ACKNOWLEDGEMENTS

For the acknowledgements I will drop the more formal, unbiased, non-anecdotic style that is so common in most research papers and theses and continue on a more personal note, for obvious reasons.

First and foremost, doing this project was an immense pleasure and a big learning experience and for me personally a big step into the critical research-world. It turned out to be a very good choice to do this project in a new environment, which perhaps piqued my interest in the matter-at-hand in new ways and challenged me to try harder to integrate into the new environment. Also, being new to the subject, I came in without much experience and thus was also more or less forced to seek help and advice of many others; this seeking of help and advice in itself was already an invaluable experience, only adding upon the obvious value of the advice itself. The opportunity to work in a lab with dedicated, creative and intelligent people and discuss the points of your project and their own work provided me with many new viewpoints and useful experiences. When I started and looked at some master theses, I had some worries if I was ever able to fill even half of such a document. Now that I was actually working on it I am sometimes wondering how I am ever going to fit everything on such a small amount of pages (and as can be noted I've had some problems with that).

Thanks to this experience I now also realize that while I learned so much during this project, it is only the beginning, and there is still a long road to go if I ever hope to make some of the valuable contributions that I have seen under construction at the VR lab. It has shown me much of what goes on in the 'research-world' and as such was a wonderful experience which taught me so much that you just can't learn just doing courses and attending lectures. However it also showed to me the use these courses had, by applying much of what I have learnt, and especially my ability to understand and analyze complex problems, on a large-scale project.

Now, without any more small random thoughts thrown in here, here are some of the people I would like to thank for helping me bring this project to a (hopefully successful) conclusion:

Konrad Tollmar – My supervisor in Lund who kept challenging my views and notes and made me look into the subject matter quite deeply. The word 'critical thinking' got a new meaning thanks to Dr. Tollmar. Being critical towards me while still stimulating me in my ideas was expertly done and a very good learning experience, especially considering his critique was usually very constructive. He also sort of gently nudged me towards other people with whom I should have a conversation about various aspects of my project and as such brought me into contact with some other people on this list.

Paul De Bra – My other supervisor from Eindhoven, who had numerous useful references, anecdotes and examples ready to liven up the material. He also kept reassuring me that it was okay that my project was not necessarily a normal one for the Computer Science faculty at home, which was always a comforting thought. Also always being critical to anything I wrote (including some silly spelling mistakes), his feedback and advice on how to better certain small issues are what put the dots on some i's and crossed some t's (in no way implying that this thesis is perfect, but it would have been worse off without his comments)

Mattias Wallergård – The professor at the lab to whom I would go with questions about the usability testing. He provided me with much help, commenting on how I should set up the test, 'testing the test' and giving me his comments, always taking as much time as was needed. The usability test was so much the better for his input. Mattias always showed a great interest in what I was doing and that proved to motivate me that much more, especially after the mention that he would definitely start up a follow-up project based on my work.

Joakim Eriksson – The 'boss' of the Virtual Reality lab and as such for me the go-to-guy on many issues I had there. He always took the time and used his impressive VR engineering knowledge to answer all of my many questions, unrelated they might sometimes be to the project. He made sure I felt comfortable by always stressing that I should have everything I need to do this project and he would do his best to provide it. Furthermore his experience with EON was invaluable and he gave me some practical tips on using the software as well as some of the hardware involved.

John Wentworth – Always a sounding board for any of my ideas or more 'philosophical' perspectives on VR, usability and anything else pretty much. John proved to be a great conversation partner always provoking new ideas and thoughts and seeing new possibilities in the smallest of things, which was also just a lot of fun!

Everyone who participated during the usability study – Everyone who helped out during this time is much obliged. It was hard to find test subjects during the vacation but these people all did it without any sort of coaxing and sacrificed some of their time for me. It sure was fun to test each and every one of them and their input will mean a lot; especially for the follow up project.

The people at home – Last but certainly not least I would like to thank the people back in the Netherlands, friends & family, who supported me while going abroad and were always interested in my ramblings about my stay in Sweden and for staying in touch, I never felt homesick thanks to that. Very special thanks to my girlfriend Lisanne who, despite obviously not wanting to see me leave for eight months, was incredibly understanding, fully supported me and made it that much easier for me to leave home for eight months to do an incredible project and have an incredible experience. Thanks for giving me so much space; I am a very lucky guy!

A. APPENDIX: CITED WORKS

Bowman, D. A., & Hodges, L. F. (1997). An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments. *Proceedings of the 1997 symposium on Interactive 3D graphics*, (p. 35).

Bowman, D. A., & Wingrave, C. A. (2001). Design and Evaluation of Menu Systems for Immersive Virtual Environments. *IEEE Virtual Reality Conference 2001*, (p. 149).

Bowman, D. A., Gabbard, J., & Hix, D. (2001). *Usability Evaluation in Virtual Environments: Classification and Comparison of Methods. Technical Report TR-01-17.* Computer Science, Virginia Tech.

Bowman, D., Koller, D., & Hodges, L. (1997). *A Methodology for the Evaluation of Travel Techniques for Immersive Virtual Environments.* Atlanta.

Bowman, D., Koller, D., & Hodges, L. (1997). *Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques.* Atlanta: Georgia Institute of Technology.

Bowman, D., Kruijf, E., LaViola, J., & Poupyrev, I. (2001). An Introduction to 3-D User Interface Design. *Presence: Teleoperators and Virtual Environments* , 96-108.

Bowman, D., Kruijff, E., LaViola, J., & Poupyrev, I. (2004). *3D User Interfaces - Theory and Practice.* Addison-Wesley.

Bryson, S. (1994). *Approaches to Successful Design and Implementation of VR Applications.* ACM SIGGRAPH.

Jegers, K., & Wiberg, C. (2003). FunTain: Design Implications for Edutainment Games. *ED-MEDIA 2003, Association for the Advancement of Computing in Education.* Charlottesville.

Kay M. Stanney. (2002). *Handbook of Virtual Environments.* New Jersey: Lawrence Erlbaum Associates, Inc., Publishers.

Kim, Y.-S., Kesavadas, T., & Paley, S. M. (2006). The Virtual Site Museum: A Multi-Purpose, Authoritative, and Functional Virtual Heritage Resource. *Presence: Teleoperators and Virtual Environments* , 245-261.

Kjeldskov, J. (2001). Combining interaction techniques and display types for virtual reality. *Proceedings of OzCHI 2001.* Edith Cowan University Press.

Mine, M. R. (1995). *Virtual Environment Interaction Techniques.* Chapel Hill: University of North Carolina.

Nelson, T. (2004). A Cosmology for a Different Computer Universe: Data Model, Mechanisms, Virtual Machine and Visualization Infrastructure. *Journal of Digital Information* , Article no. 298, alsa available online at http://jodi.tamu.edu/Articles/v05/i01/Nelson/.

Nielsen, J. a. (1993). A mathematical model of the finding of usability problems. *Proceedings of ACM INTERCHI'93 Conference*, (pp. 206-213). Amsterdam.

Parés, N., & Parés, R. (2006). Subjectiveness, Towards a Model for a Virtual Reality Experience: The Virtual Subjectiveness. *Presence: Teleoperators and Virtual Environments* , 524-538.

Preece, J., Rogers, Y., & Sharp, H. (2002). *Interaction Design – Beyond human-computer interaction.* John Wiley & Sons, Inc.

Rutledge, L., Aroyo, L., & Stash, N. (2006). Determining User Interests About Museum Collections. *The 15th International Conference on World Wide Web (WWW'06) .* Edinburgh.

Schroeder, R., Heldal, I., & Tromp, J. (2006). The Usability of Collaborative Virtual Environments and Methods for the Analysis of Interaction. *Presence: Teleoperators and Virtual Environments* , 655-667.

Silver, M. (2005). *Exploring Interface Design.* Thomson Delmar Learning.

Steed, A., & Parker, C. (2005). Evaluation Effectiveness of Interaction Techniques across Immersive Virtual Environmental Systems. *Presence* , 511-527.

Tromp, J. G., Steed, A., & Wilson, J. R. (2003). Systematic Usability Evaluation and Design Issues for Collaborative Virtual Environments. *Presence: Teleoperators and Virtual Environments* , 241-267.

Turk, M. (1998). *Moving from GUIs to PUIs.* Redmond: Microsoft Corporation.

Verheijen, B. (2004). *An Experiment on Two-Handed Interaction in the Personal Space Station.* Eindhoven: Eindhoven University of Technology.

Vermeulen, G. (2008). 3D Input in een Virtuele Omgeving. Genk: Media & Design Academie.

Wallergård, M. (2007). Initial Usability Testing of Navigation and Interaction Methods in Virtual Environments: Developing Usable Interfaces for Brain Injury Rehabilitation. *Presence* , 16-44.

Wang, Y. (2007). User-Centered Design for Personalized Access to Cultural Heritage. *11th International Conference on User Modeling.* Greece.

Wiberg, C. (2003). *A Measure of Fun. Extending the scope of web usability.* Umeå: Department of Informatics. Umeå University.

Youngblut, C., Johnson, R., Nash, S., Wienclaw, R., & Will, C. (1996). *Review of virtual environment interface technology.* Institute for Defense Analysis.

B. APPENDIX: USER STUDY – QUESTIONNAIRE

This is the questionnaire as it was presented to users who completed the test with the on-screen hints. Users without these hints did not get to see the questions about their influence on the immersion, located in the immersion section of this questionnaire.

Post test questionnaire

User:

Aesthetic

Would you have liked a bigger room of exhibits? If yes: How much bigger?	Yes - No Slightly bigger – Twice as big – Much bigger				
Would you have liked more rooms? If yes: How many more?	Yes – No O 1	O 2-4	O 5-10	O 10-20	O >20
Was the level of detail on exhibits adequate?	O Way too little	O	O Exactly enough	O	O Way too much

Comfort

Would you prefer to sit or to stand?	SpaceBall: Sit – Stand WiiMote: Sit – Stand				
Are you feeling any strain after using the device	SpaceBall: Yes – No WiiMote: Yes – No				
If you are feeling strain, is it more than using a normal computer?	O Much Less	O	O The Same	O	O Much More

Ease of Learning

Which device did you find easier to understand?	SpaceBall – WiiMote					
Please rank these tasks in order of difficulty (1 = easiest, 4 = hardest)	SB WM 	-Walking around -Looking around -Operating the selection wheel -Selecting an object				
How easy was it to learn: *How to access the educational game?*	O I had no idea	O	O Took me a while	O	O Understood it at once	
How to enter the historical simulation?	O I had no idea	O	O Took me a while	O	O Understood it at once	
How to find related content?	O I had no idea	O	O Took me a while	O	O Understood it at once	
How to stroll around the museum and the simulation?	O I had no idea	O	O Took me a while	O	O Understood it at once	

Interactivity

Did you miss the ability to manipulate exhibits?	Yes – No				
How natural would you say the device worked: *SpaceBall*	O Completely not as expected	O	O So-So	O	O Completely Natural
WiiMote	O Completely not as expected	O	O So-So	O	O Completely Natural

Immersion

How much would you say this felt like a real museum visit?	O Not even close O O Somewhat the Same O O Completely the same
Was the historical simulation believable?	Yes – No
Did you find the sense of immersion important when	Walking through the museum: Yes – No Exploring related information: Yes – No In the historical simulation: Yes - No
Did the onscreen hints have an influence on your sense of immersion If yes, how much	Yes – No O A little bit less O O Quite a bit O O Made it completely fake

How much did you appreciate the following features?

Feature	Appreciation				
Walking around and looking at the same time	O Didn't care for it	O	O It was okay	O	O Very good
Browsing related content	O Didn't care for it	O	O It was okay	O	O Very good
Having small questions about the museum	O Didn't care for it	O	O It was okay	O	O Very good
The historical simulation	O Didn't care for it	O	O It was okay	O	O Very good
The ZoomBack travel technique (automatic movement towards exhibits)	O Didn't care for it	O	O It was okay	O	O Very good

C. APPENDIX: USER STUDY – INTERVIEW

This is the interview as it was presented to users participating in the user study with on screen hints. The question about the onscreen hints was of course left out for those who did a test without the on screen hints. Note how the last question is actually answerable with just a yes or a no, however this question is important enough (one could say the central question to any usability study) to warrant an answer and the explanation of this answer could be important.

User:

Age:

Sex:

Starting device:

User's technical background concerning 3D navigation and the WiiMote:

Expectations

How do you expect to be able to walk around?

What do you imagine when I say virtual game in a museum?

If I talk about a big historical simulation what sort of environment do you expect in the context of this Kulturen museum (Medieval themed)?

If I talk about 'related content' to an exhibit what sort of content would you expect?

Would the interface (so not the device) matter to you when you think about a virtual museum and what are your expectations?

Post test interview questions

What was the thing that appealed to you most in this simulation you just did?

And what appealed to you the least?

How did you like the two devices?

What did you think of the educational game?

What about the big simulation?

How would you compare this experience to a 'true immersive experience' (e.g. with stereoscopic glasses)?

Is there anything regarding the interactiveness which you would like to see extended?

What about new possibilities in this museum?

[Discussion on the questionnaire]

How did this experience measure up to a real museum with regards to education and fun?

Do you have any comments on the onscreen hints; and lack thereof in the historical simulation?

And finally, would you actually use an application like this if you came across it?

D. APPENDIX: FULL TABLE OF TASK PERFORMANCE RESULTS

This is the full table of results for test subject 1 through 15 in random order

	T_1	T_2	T_3	T_4	T_5	T_6	T_7
Age	46	20	29	25	27	22	24
Sex	F	M	M	M	F	F	M
WM Exp	Y	Y	N	N	N	N	N
SB Exp	N	N	N	N	N	N	N
3D exp	Y	Y	N	Y	N	N	Y
Order	1) SB 2) WM	1) SB 2) WM	1) WM 2) SB	1) WM 2) SB	1)SB 2) WM	1) SB 2) WM	1) WM 2) SB
OSH	N	N	N	N	N	Y	Y
	SpaceBall	**SpaceBall**	**WiiMote**	**WiiMote**	**SpaceBall**	**SpaceBall**	**WiiMote**
Task 1							
Time	60	151	112	63	239	121	61
Errors	2	2	6	3	10	2	4
Faults	0	1	0	1	3	0	0
Task 2							
Time	31	17	95	78	62	55	15
Errors	0	0	4	6	3	1	1
Faults	1	0	2	0	0	1	0
Task 3							
Time	52	46	103	32	58	46	34
Errors	0	2	4	0	1	0	0
Faults	1	1	1	1	0	1	1

	WiiMote	WiiMote	SpaceBall	SpaceBall	WiiMote	WiiMote	SpaceBall
Task 1							
Time	63	26	103	104	72	38	71
Errors	2	0	6	2	4	0	5
Faults	0	0	1	0	0	0	0
Task 2							
Time	50	38	41	18	93	35	14
Errors	3	1	2	0	4	2	0
Faults	0	0	0	0	0	0	0
Task 3							
Time	31	17	30	20	42	17	39
Errors	0	0	0	0	0	0	0
Faults	0	0	0	0	0	0	0

Age	**T_8**	**T_9**	**T_{10}**	**T_{11}**	**T_{12}**	**T_{13}**	**T_{14}**	**T_{15}**
Sex	21	22	21	54	24	35	31	31
WM Exp	M	M	F	M	M	M	F	M
SB Exp	N	N	N	N	Y	Y	N	N
3D exp	N	N	N	N	N	N	N	Y
Order	Y	N	Y	Y	Y	Y	N	Y
OSH	1) SB 2) WM	1) WM 2) SB	1) WM 2) SB	1) WM 2) SB	1) SB 2) WM	1)SB 2) WM	1) WM 2) SB	1) WM 2) SB
	Y	N	Y	Y	Y	Y	N	Y
Task 1	**SpaceBall**	**WiiMote**	**WiiMote**	**WiiMote**	**SpaceBall**	**SpaceBall**	**WiiMote**	**WiiMote**

Time								
Errors	193	117	53	167	42	73	156	46
Faults	7	5	3	9	1	2	5	1
Task 2	0	1	1	2	0	1	3	2
Time								
Errors	25	111	161	101	21	29	196	65
Faults	1	5	4	3	0	0	7	3
Task 3	0	2	3	2	0	1	3	0
Time								
Errors	33	35	55	81	51	30	63	39
Faults	1	0	0	0	0	0	0	0
	0	0	1	2	2	0	1	0
Task 1	**WiiMote**	**SpaceBall**	**SpaceBall**	**SpaceBall**	**WiiMote**	**WiiMote**	**SpaceBall**	**SpaceBall**
Time								
Errors	127	91	65	156	27	54	240	44
Faults	3	4	4	4	0	1	14	2
Task 2	0	1	0	1	0	0	0	0
Time								
Errors	75	18	16	25	27	54	22	25
Faults	5	0	0	0	2	3	0	1
Task 3	0	0	0	0	0	0	0	0
Time								
Errors	45	21	22	41	51	45	44	24
Faults	1	0	0	2	0	0	0	0
	0	0	0	0	0	0	0	0

E. APPENDIX: FULL TABLE OF QUESTIONNAIRE RESULTS

These are the questionnaire results for test subjects 1 through 15 in random order.

	T_1	T_2	T_3	T_4	T_5	T_6	T_7
Aesthetic							
Bigger Room	No	Yes	No	Yes	No	No	Yes
How Much		3.00		2.00			2.00
More Rooms?	Yes	Yes	No	Yes	Yes	Yes	Yes
How Many	3.00	Depends		2.00	3.00	2.00	4.00
LOD Adequate?	1.00	1.00	2.00	3.00	2.00	2.00	1.00
Comfort							
Sit/Stand(SB)	Sit	Sit/Stand	Sit	Sit	Sit	Sit	Sit
Sit/Stand(WM)	Sit/Stand	Stand	Stand	Stand	Sit	Stand	Stand
Strain(SB)	No	No	No	Yes	Yes	No	No
Strain(WM)	No	No	Yes	No	No	No	No
If Strain-More?			4.00	5.00	2.00		
Learning							
Easier device	WiiMote	SpaceBall	SpaceBall	WiiMote	WiiMote	WiiMote	WiiMote
Order of Difficulty							
SB-Walking Around	3.00	4.00	1.00	4.00	4.00	4.00	4.00
Looking Around	2.00	3.00	3.00	3.00	3.00	3.00	3.00
Operating wheel	4.00	2.00	4.00	2.00	1.00	2.00	2.00
Selecting	1.00	1.00	2.00	1.00	2.00	1.00	1.00
WM-Walking Around	3.00	4.00	1.00	1.00	2.00	2.00	2.00
Looking Around	2.00	3.00	3.00	4.00	4.00	1.00	4.00
Operating wheel	4.00	2.00	4.00	3.00	3.00	4.00	3.00
Selecting	1.00	1.00	2.00	2.00	1.00	3.00	1.00
How easy was it to:							
Learn Edu. Game	5.00	5.00	3.00	5.00	4.00	5.00	4.00
Enter Simulation	5.00	4.00	5.00	5.00	4.00	4.00	5.00
Find related content	4.00	5.00	3.00	5.00	4.00	5.00	5.00
Stroll around	3.00	4.00	4.00	5.00	3.00	5.00	3.00
Interactivity							
Miss manipulate	No	Yes	Yes	Yes	No	No	Yes
Natural (SB)	3.00	4.00	4.00	1.00	1.00	1.00	3.00
Natural (WM)	3.00	1.00	2.00	5.00	3.00	4.00	4.00
Immersion							
Real Museum Visit?	1.00	3.00	2.00	1.00	1.00	4.00	2.00
Simulation Believable?	No	No	Yes	Yes	Yes	Yes	No
Immersion important when							
Walking through	No	Yes	No	Yes	Yes	Yes	Yes

Exploring information	No	No	No	Yes	No	No	Yes
In simulation	No	Yes	Yes	Yes	Yes	Yes	Yes
OSH influence?						No	No
How much?							
Features							
Browsing content	3.00	5.00	3.00	4.00	5.00	5.00	5.00
Small quiz	2.00	5.00	3.00	4.00	4.00	3.00	2.00
Simulation	2.00	5.00	4.00	5.00	5.00	5.00	3.00
ZoomBack	5.00	5.00	4.00	5.00	3.00	5.00	5.00

	T_8	T_9	T_{10}	T_{11}	T_{12}	T_{13}	T_{14}	T_{15}
Aesthetic								
Bigger Room	No	Yes	No	No	Yes	No	Yes	Yes
How Much		2.00			3.00		2.00	2.00
More Rooms?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
How Many	3.00	4.00	2.00	3.00	2.00	3.00	3.00	
LOD Adequate?	3.00	2.00	2.00	1.00	3.00	1.00	1.00	1.00
Comfort								
Sit/Stand(SB)	Sit	Stand	Sit	Stand	Sit	Sit	Sit	Sit
Sit/Stand(WM)	Stand	Stand	Sit/Stand	Stand	Sit	Stand	Sit	Sit
Strain(SB)	No	No	No	No	No	No	Yes	No
Strain(WM)	No	No	No	No	No	No	No	No
If Strain-More?							5.00	
Learning								
Easier device	WiiMote	SpaceBall	SpaceBall	SpaceBall	SpaceBall	SpaceBall	WiiMote	WiiMote
Order of Difficulty								
SB-Walking Around	4.00	4.00	3.00	4.00	4.00	2.00	4.00	4.00
Looking Around	3.00	1.00	2.00	3.00	1.00	3.00	3.00	1.00
Operating wheel	2.00	3.00	4.00	2.00	2.00	1.00	2.00	2.00
Selecting	1.00	2.00	1.00	1.00	3.00	4.00	1.00	3.00
WM-Walking Around	1.00	2.00	2.00	2.00	1.00	1.00	3.00	1.00
Looking Around	2.00	4.00	3.00	1.00	4.00	4.00	2.00	4.00
Operating wheel	4.00	3.00	4.00	3.00	3.00	3.00	4.00	3.00
Selecting	3.00	1.00	1.00	4.00	2.00	2.00	1.00	2.00

How easy was it to:								
Learn Edu. Game	5.00	4.00	1.00	3.00	5.00	4.00	3.00	5.00
Enter Simulation	5.00	3.00	5.00	3.00	5.00	5.00	3.00	3.00
Find related content	5.00	5.00	5.00	5.00	5.00	4.00	3.00	4.00
Stroll around	3.00	4.00	3.00	4.00	5.00	5.00	3.00	4.00
Interactivity								
Miss manipulate	No	Yes	No	No	No	No	Yes	Yes
Natural (SB)	2.00	4.00	4.00	4.00	5.00	4.00	2.00	4.00
Natural (WM)	4.00	3.00	4.00	3.00	3.00	2.50	3.00	3.00
Immersion								
Real Museum Visit?	3.00	4.00	3.00	3.50	3.00	2.00	1.00	3.00
Simulation Believable?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Immersion important when								
Walking through	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Exploring information	Yes	No	No	Yes	No	Yes	Yes	Yes
In simulation	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
OSH influence?	No		No	No	No	No		Yes
How much?								3.00
Features								
Browsing content	5.00	5.00	4.00	4.00	5.00	5.00	5.00	5.00
Small quiz	4.00	4.00	5.00	5.00	5.00	5.00	5.00	3.00
Simulation	4.00	5.00	5.00	4.00	4.00	5.00	3.00	4.00
ZoomBack	4.00	5.00	5.00	5.00	5.00	5.00	5.00	5.00